EXAMTOPICS

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



CERTIFICATION TEST

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

Question #1 *Topic 1*

Your company has decided to make a major revision of their API in order to create better experiences for their developers. They need to keep the old version of the API available and deployable, while allowing new customers and testers to try out the new API. They want to keep the same SSL and DNS records in place to serve both APIs.
What should they do?

A. Configure a new load balancer for the new version of the API

B. Reconfigure old clients to use a new endpoint for the new API

C. Have the old API forward traffic to the new API based on the path

D. Use separate backend pools for each API path behind the load balancer

Suggested Answer: *D*
Official Answer: D

Community vote distribution
D (100%)

shandy Highly Voted 1 year, 11 months ago
D is the answer because HTTP(S) load balancer can direct traffic reaching a single IP to different backends based on the incoming URL. A is not correct because configuring a new load balancer would require a new or different SSL and DNS records which conflicts with the requirements to keep the same SSL and DNS records. B is not correct because it goes against the requirements. The company wants to keep the old API available while new customers and testers try the new API. C is not correct because it is not a requirement to decommission the implementation behind the old API. Moreover, it introduces unnecessary risk in case bugs or incompatibilities are discovered in the new API.D is right
upvoted 116 times

shandy Highly Voted 1 year, 11 months ago
D is the answer because HTTP(S) load balancer can direct traffic reaching a single IP to different backends based on the incoming URL. A is not correct because configuring a new load balancer would require a new or different SSL and DNS records which conflicts with the requirements to keep the same SSL and DNS records. B is not correct because it goes against the requirements. The company wants to keep the old API available while new customers and testers try the new API. C is not correct because it is not a requirement to decommission the implementation behind the old API. Moreover, it introduces unnecessary risk in case bugs or incompatibilities are discovered in the new API.
upvoted 116 times

AzureDP900 3 years, 10 months ago
D is right
upvoted 4 times

AWS56 6 years, 9 months ago
agreed, The answer is D
upvoted 20 times

Wayne90040 Most Recent 3 months, 3 weeks ago
D is the answer
upvoted 1 times

Dinya_jui 7 months, 2 weeks ago
As we need to keep the old api's and use the same ssl and dns, using separate backends would be the correct approach
upvoted 1 times

edudope 8 months, 1 week ago
None of the new Oct 30,2025 exam questions is here, does anyone know when they will update it?I asked support - they are not even aware, stating they updated the exam in December. So do not expect any updates soon.... Total waste of money!
upvoted 3 times

edudope 8 months, 1 week ago
Selected Answer: D
None of the new Oct 30,2025 exam questions is here, does anyone know when they will update it?
upvoted 3 times

ToiToi 7 months, 4 weeks ago

I asked support - they are not even aware, stating they updated the exam in December. So do not expect any updates soon.... Total waste of money!
upvoted 4 times

Arnabbabu 8 months, 4 weeks ago
Please note that None of the questions in This Question set appear in the exam. The claim is false
upvoted 9 times

tuxtion 11 months ago
agreed, The answer is D
upvoted 1 times

pjones301 1 year ago
Ans: D. To keep the old end points we need to create a new backend app for new endpoint path which will route through the load balancer. In the DNS Zone we need to put the same A record (ALB IP) with the new API subdomain.
upvoted 1 times

mgrmek 1 year ago
D is right
upvoted 1 times

Minkio 1 year, 2 months ago
D is right
upvoted 1 times

WakandaF 1 year, 5 months ago
I need to perform the Google Architect Renew. There's a Google release of one Architect certification called "Professional Cloud Architect Renewal." I would like to know if the Google Architect certification available on this site covers these renewal questions. Thanks in advance.
https://services.google.com/fh/files/misc/pca_renewal_exam_overview_and_faqs.pdf CQ
upvoted 6 times

j_vish 1 year, 5 months ago
D is correct
upvoted 1 times

biplabid 1 year, 6 months ago
Answer is D. Using a HTTPS load balancer for different backend makes sense as we dont want to disturb the old API setup and also want to test new API
upvoted 1 times

Test88888 1 year, 6 months ago
Can anybody tell if these here question are still in the current exam?
upvoted 2 times

XiaobinJiang 1 year, 7 months ago
D is right!
upvoted 1 times

FI22 1 year, 8 months ago
D is good for keep the same records.
upvoted 1 times

Ghorbel 1 year, 8 months ago
Using separate backend pools for each API version behind a single load balancer fulfills all requirements: Single SSL and DNS Endpoint: Both API versions share the same SSL certificate and DNS records, so users and developers don't need to reconfigure their endpoints. The load balancer handles routing transparently. Separation of APIs: Different backend pools allow for clear segregation between the old and new API versions. Each API can be deployed, updated, and scaled independently. Traffic Management: The load balancer can route traffic based on request paths (e.g., /v1 for the old API and /v2 for the new API). This allows smooth coexistence of the APIs while developers gradually adopt the new version.
upvoted 1 times

Question #2 *Topic 1*

Your company plans to migrate a multi-petabyte data set to the cloud. The data set must be available 24hrs a day. Your business analysts have experience only with using a SQL interface.
How should you store the data to optimize it for ease of analysis?

A. Load data into Google BigQuery

B. Insert data into Google Cloud SQL

C. Put flat files into Google Cloud Storage

D. Stream data into Google Cloud Datastore

Suggested Answer: *A*
BigQuery is Google's serverless, highly scalable, low cost enterprise data warehouse designed to make all your data analysts productive. Because there is no infrastructure to manage, you can focus on analyzing data to find meaningful insights using familiar SQL and you don't need a database administrator. BigQuery enables you to analyze all your data by creating a logical data warehouse over managed, columnar storage as well as data from object storage, and spreadsheets. Reference: https://cloud.google.com/bigquery/

Community vote distribution
A (97%) (3

Eroc Highly Voted 1 year, 11 months ago
This question could go either way for A or B. But Big Query was designed with this in mind, according to numerous Google presentation and videos. Cloud Datastore is a NoSQL database (https://cloud.google.com/datastore/docs/concepts/overview) Cloud Storage does not have an SQL interface. The previous two sentences eliminate options C and D. So I'd pick "A".IMHO, it should be A only. The reason is that they want to perform analysis on the data and BigQuery excels in that over Cloud SQL. You can run SQL queries in both but I BigQuery has better analytical tools. It can do ad-hoc analysis like Cloud SQL using Cloud Standard SQL and it can do geo-spatial and ML analysis via its Cloud Standard SQL interface.A is okCloud SQL does not scale to that magnitude also Cloud SQL is not meant for OLAP Answer is BigQueryB is not correct because Cloud SQL storage limit doesn't fit the requirement.
upvoted 45 times

Eroc Highly Voted 1 year, 11 months ago
This question could go either way for A or B. But Big Query was designed with this in mind, according to numerous Google presentation and videos. Cloud Datastore is a NoSQL database (https://cloud.google.com/datastore/docs/concepts/overview) Cloud Storage does not have an SQL interface. The previous two sentences eliminate options C and D. So I'd pick "A".
upvoted 44 times

0xE8D4A51000 3 years, 10 months ago
IMHO, it should be A only. The reason is that they want to perform analysis on the data and BigQuery excels in that over Cloud SQL. You can run SQL queries in both but I BigQuery has better analytical tools. It can do ad-hoc analysis like Cloud SQL using Cloud Standard SQL and it can do geo-spatial and ML analysis via its Cloud Standard SQL interface.
upvoted 2 times

tartar 6 years ago
A is ok
upvoted 16 times

zr79 3 years, 10 months ago
Cloud SQL does not scale to that magnitude also Cloud SQL is not meant for OLAP Answer is BigQuery
upvoted 5 times

kinghin 4 years, 7 months ago
B is not correct because Cloud SQL storage limit doesn't fit the requirement.
upvoted 15 times

clouddude 6 years, 3 months ago
I'll go with A because BQ (and BT) are usually meant for analytics. B isn't correct because Cloud SQL does not scale to that volume. C isn't correct because Cloud Storage does not provide a standard SQL mechanism. D could be right but it sounds off because of the analytics requirement.
upvoted 15 times

pjones301 Most Recent 1 year ago

GBQ supports multi PB of records, CloudSQL won't.

upvoted 1 times

mgrmek 1 year ago

A is ok

upvoted 1 times

Sib09 1 year, 5 months ago

Cloud SQL supports TB not PB of data

upvoted 3 times

RTATAVAR 1 year, 5 months ago

optimize the ease of Analysis ... Any data that can be analyzed is using Big Query So Option A will suit the requirement

upvoted 1 times

lks68 1 year, 6 months ago

CloudSQL cannot handle petabytes but tens of terabytes only.

upvoted 1 times

szagarella 1 year, 6 months ago

Cloud SQL is limited to 30TB of data. A is correct.

upvoted 1 times

hpf97 1 year, 6 months ago

A or B but as it said it is for an analyst, A is better

upvoted 1 times

XiaobinJiang 1 year, 7 months ago

I'll go A as BQ has compatible usage of SQL

upvoted 1 times

devenderpraksh 1 year, 9 months ago

Answer: A. Load data into Google BigQuery This is the correct choice because it's the only option that meets all requirements: Can handle multi-petabyte scale Provides the required SQL interface for analysts Ensures 24/7 availability Optimized for analytical queries Serverless and automatically scales

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

The best answer here is A. Load data into Google BigQuery. Here's why: 1. Designed for large datasets: BigQuery is a serverless, highly scalable, and cost-effective multicloud data warehouse designed specifically for analyzing massive datasets. Petabyte-scale data is exactly what it excels at. 2. SQL interface: Your analysts are already familiar with SQL, and BigQuery uses standard SQL, making the transition easy and minimizing the learning curve. 3. High availability: BigQuery offers high availability with built-in redundancy and replication. 4. Performance: BigQuery is optimized for analytical queries and can handle complex queries across massive datasets very efficiently.

upvoted 1 times

dev_evening 1 year, 10 months ago

A, BigQuery is more suitable for analysis. While Cloud SQL can work, it's more generic

upvoted 1 times

amxexam 1 year, 11 months ago

Let's go with option elimination A. Load data into Google BigQuery >>Big Query = Analytic + SQL (Ease of using SQL) Storage hence the solution B. Insert data into Google Cloud SQL >> Yes you can SQL query with your own applicaiton console compared to BigQuery SQL console, and 24 hrs avalablity but you won't have 1-2 sec response on petabytes of data, as you can do in GCP BigQuery partitioned and clustered tables. C. Put flat files into Google Cloud Storage >>The requirement is for analytics and SQL querying of data. You can store it in the flat file but will need to use GCP BigQuery to do that D. Stream data into Google Cloud Datastore >> Only dealing with storage problems does not address analytics and SQL querying

upvoted 2 times

amxexam 1 year, 11 months ago

Let's go with option elimination A. Load data into Google BigQuery >>Big Query = Analytic + SQL (Ease of using SQL) Storage hence the solution B. Insert data into Google Cloud SQL >> Yes you can SQL query with your own applicaiton console compared to BigQuery SQL console, and 24 hrs avalablity but you won't have 1-2 sec response on petabytes of data, as you can do in GCP BigQuery partitioned and clustered tables. C. Put flat files into Google Cloud Storage >>The requirement is for analytics and SQL querying of data. You can store it in the flat file but will need to use GCP BigQuery to do that D. Stream data into Google Cloud Datastore >> Only dealing with storage problems does not address analytics and SQL queryingHence Option A

upvoted 2 times

amxexam 4 years, 11 months ago

Hence Option A

upvoted 2 times

i_am_robot 1 year, 11 months ago

A. Load data into Google BigQuery BigQuery is a fully managed, cloud-native data warehousing solution that makes it easy to analyze large and complex datasets. It is optimized for analyzing large amounts of data quickly, and can handle petabyte-scale datasets with ease. It also has a SQL-like interface that is familiar to business analysts, making it easy for them to query and analyze the data. Additionally, BigQuery is highly scalable and can handle high query concurrency, making it a good choice for storing data that must be available 24/7. Option B, inserting data into Google Cloud SQL, is not a good choice for a multi-petabyte dataset because Cloud SQL is not designed to handle such large volumes of data. Option C, putting flat files into Cloud Storage, is also not a good choice because it is not optimized for querying and analyzing data. Option D, streaming data into Cloud Datastore, is not a good choice because Cloud Datastore is a NoSQL database and does not have a SQL-like interface.

upvoted 2 times

omermahgoub 1 year, 11 months ago

A. Load data into Google BigQuery To optimize the storage of the multi-petabyte data set for ease of analysis by business analysts who have experience only with using a SQL interface, you should load the data into Google BigQuery. BigQuery is a fully-managed, cloud-native data warehouse that allows you to perform fast SQL queries on large amounts of data. By loading the data into BigQuery, you can provide your business analysts with a familiar SQL interface for querying the data, making it easier for them to analyze the data set. Other options, such as inserting data into Google Cloud SQL, putting flat files into Google Cloud Storage, or streaming data into Google Cloud Datastore, may not provide the necessary SQL interface or query performance for efficient analysis of the data set.

upvoted 2 times

juanlopezcervero 2 years, 3 months ago

A is correct

upvoted 1 times

Question #3 *Topic 1*

The operations manager asks you for a list of recommended practices that she should consider when migrating a J2EE application to the cloud. Which three practices should you recommend? (Choose three.)

A. Port the application code to run on Google App Engine

B. Integrate Cloud Dataflow into the application to capture real-time metrics

C. Instrument the application with a monitoring tool like Stackdriver Debugger

D. Select an automation framework to reliably provision the cloud infrastructure

E. Deploy a continuous integration tool with automated testing in a staging environment

F. Migrate from MySQL to a managed NoSQL database like Google Cloud Datastore or Bigtable

Suggested Answer: *A, D, E*

References: https://cloud.google.com/appengine/docs/standard/java/tools/uploadinganapp
https://cloud.google.com/appengine/docs/standard/java/building-app/cloud-sql

Community vote distribution



NapoleonBorntoparty Highly Voted 1 year, 11 months ago

This is talking about the APPLICATION not the infrastructure, therefore I believe we should focus on the APP-side of things: 1. port the app to app engine for content delivery 2. add monitoring for troubleshooting 3. use a CI/CD workflow for continuous delivery w/testing for a stable application

so, for me: A, C and E should be the answers

upvoted 81 times

NapoleonBorntoparty Highly Voted 1 year, 11 months ago

This is talking about the APPLICATION not the infrastructure, therefore I believe we should focus on the APP-side of things: 1. port the app to app engine for content delivery 2. add monitoring for troubleshooting 3. use a CI/CD workflow for continuous delivery w/testing for a stable application so, for me: A, C and E should be the answersthe person who asking you recommendation is operation manager, it can be related to infrastructureA is not... This is a refactoring/re-platforming step, not a universal practice. While beneficial, the initial migration might focus on lift-and-shift to Compute Engine or GKE first (less invasive). This is a strategic choice, not a recommended practice for all migrations.

upvoted 81 times

segkhachat 3 years, 5 months ago

the person who asking you recommendation is operation manager, it can be related to infrastructure

upvoted 5 times

ast3citos 10 months, 1 week ago

A is not... This is a refactoring/re-platforming step, not a universal practice. While beneficial, the initial migration might focus on lift-and-shift to Compute Engine or GKE first (less invasive). This is a strategic choice, not a recommended practice for all migrations.

upvoted 2 times

amxexam 4 years, 12 months ago

Let's go with option elimination A. Port the application code to run on Google App Engine >> PaaS serverless managed service, so all my infra provisioning is taken care by GCP. B. Integrate Cloud Dataflow into the application to capture real-time metrics >> Good to have C. Instrument the application with a monitoring tool like Stackdriver Debugger >> Is a must for debugging issues and monitoring application logs this is now GCP Cloud monitoring and logging. D. Select an automation framework to reliably provision the cloud infrastructure >> App Engine is a PaaS so the infrastructure is taken care of by App Engine, I would select this if I have not selected A, hence will eliminate this option for now E. Deploy a continuous integration tool with automated testing in a staging environment >> Good to have F. Migrate from MySQL to a managed NoSQL database like Google Cloud Datastore or Bigtable >> There is no requirement for DB enhancement hence will elimination this option A and C are must-have B and E are Good to have, but E has more importance than Big Hence will go with ACE

upvoted 17 times

exam_prepper_1234 Most Recent 4 months, 1 week ago

A is not common practice, and we don't know anything about it application. Without that information, we should not suggest using the app engine.

upvoted 1 times

ch_rehan 6 months, 3 weeks ago

Monitoring is essential to observe performance and failures after migration. Infrastructure automation ensures repeatable and reliable cloud environments. CI with automated testing minimizes deployment risk and supports safe releases.

upvoted 1 times

mgrmek 1 year ago

c d e should be

upvoted 1 times

viniciusrp2 1 year, 1 month ago

The question dont mention infrastructure. In the perspective of a APPLICATION, chose A, C, E.

upvoted 1 times

mogita 1 year, 2 months ago

Regarding A, it's not necessarily a must-have as porting application has nothing to do with operations.

upvoted 1 times

rinkujacob 1 year, 2 months ago

A, C, E are the options to recommend, Option F is not directly related to migration. Moving to cloud you should focus in PaaS rather than managing own infra of VMs, so option D is not needed.

upvoted 1 times

mahi_h 1 year, 3 months ago

D and E is confirmed. Choosing option C over A because the ask is from operations manager (i.e. Infra team)

upvoted 1 times

examxjh 1 year, 4 months ago

There is no requirement stated for B and F, so we left A, C, D and E A is important because the question focuses on an app. C is definitely needed when it comes to application development since D will be covered by Google App Engine in A, E is more useful if we are talking about best practice. so, A,C,E

upvoted 1 times

Syruis 1 year, 5 months ago

C is not the good answer because Stack Debugger as been abandonned since 2023. So it's ADE for the others reasons ...

upvoted 1 times

apb98 1 year, 5 months ago

CDE for me. C, because it’s a recommended practice to monitor the application after migration. Stackdriver Debugger (now part of the Cloud Operations suite) is useful to understand the application’s behavior without interrupting it. D, because using tools like Terraform, Deployment Manager, or similar is a key best practice when migrating to the cloud. E, because automating testing and integration is a core part of any modern cloud migration strategy.

upvoted 2 times

RTATAVAR 1 year, 5 months ago

My answer is C,D,E because choose the options which does not mention specifically about a particular service

upvoted 1 times

PatnalaJayaram 1 year, 5 months ago

I feel acd are correct

upvoted 1 times

hpf97 1 year, 6 months ago

A : GCP managed infrastructure to make JEE application work, easier than on a VM with server appli to configure C : Mandatory in order to operate the application (maybe not the debugger, but the monitoring) E : CI/CD for better application quality / security; so the operator would appreciate

upvoted 1 times

LeoSantos121212121212121 1 year, 7 months ago

I chose these answers based on what would make the application easier to get deployed to gcp.

upvoted 1 times

Ishu_awsguy 1 year, 7 months ago

I would say A C D . E has no relevance to the question , creating a CI CD automation testing pipeline that too in staging , yes is a good thing but it is not relavant I feel. Similary doing Data analysis is also irrelevant . hence B & E go out for the same reason .

upvoted 2 times

Question #4 *Topic 1*

A news feed web service has the following code running on Google App Engine. During peak load, users report that they can see news articles they already viewed.
What is the most likely cause of this problem?

```
import news
from flask import Flask, redirect, request
from flask.ext.api import status
from google.appengine.api import users

app = Flask(_name_)
sessions = {}

@app.route("/")
def homepage():
      user = users.get_current_user()
      if not user:
            return "Invalid login",
status.HTTP_401_UNAUTHORIZED

      if user not in sessions:
            sessions[user] = {"viewed": []}

      news_articles = news.get_new_news (user, sessions [user]
["viewed"])
      sessions [user] ["viewed"] +- [n["id"] for n
in news_articles]

      return news.render(news_articles)

if _name_ == "_main_":
      app.run()
```

A. The session variable is local to just a single instance

B. The session variable is being overwritten in Cloud Datastore

C. The URL of the API needs to be modified to prevent caching

D. The HTTP Expires header needs to be set to -1 stop caching

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (100%)

jackdbd Highly Voted 1 year, 11 months ago
It's A. AppEngine spins up new containers automatically according to the load. During peak traffic, HTTP requests originated by the same user could be served by different containers. Given that the variable `sessions` is recreated for each container, it might store different data. The problem here is that this Flask app is stateful. The `sessions` variable is the state of this app. And stateful variables in AppEngine / Cloud Run / Cloud Functions are problematic. A solution would be to store the session in some database (e.g. Firestore, Memorystore) and retrieve it from there. This way the app would fetch the session from a single place and would be stateless.
upvoted 128 times

jackdbd Highly Voted 1 year, 11 months ago
It's A. AppEngine spins up new containers automatically according to the load. During peak traffic, HTTP requests originated by the same user could be served by different containers. Given that the variable `sessions` is recreated for each container, it might store different data. The problem here is that this Flask app is stateful. The `sessions` variable is the state of this app. And stateful variables in AppEngine / Cloud Run / Cloud Functions are problematic. A solution would be to store the session in some database (e.g. Firestore, Memorystore) and retrieve it from there. This way the app would fetch the session from a single place and would be stateless.Very well stated, jack. I just wanted to point, GAE is a webserver platform anyway, so making application stateless or stateful is up to the developer and has nothing to do with GAE. The issue is about session consistency. GAE spin

new container if there's a need, and based on the code, the session is stored locally, this means, there's no consistency between container, and there's no grantee that the same container might serve the same user. Thank you Jack, very good explanation

upvoted 128 times

omermahgoub 3 years, 8 months ago

Very well stated, jack. I just wanted to point, GAE is a webserver platform anyway, so making application stateless or stateful is up to the developer and has nothing to do with GAE. The issue is about session consistency. GAE spin new container if there's a need, and based on the code, the session is stored locally, this means, there's no consistency between container, and there's no grantee that the same container might serve the same user. Thank you Jack, very good explanation

upvoted 16 times

JoeShmoe 6 years, 9 months ago

A is correct

upvoted 29 times

Wayne90040 Most Recent 3 months, 3 weeks ago

A is the answer.

upvoted 1 times

ybenck 1 year ago

A is good answer

upvoted 1 times

mgrmek 1 year ago

a is ok

upvoted 1 times

hpf97 1 year, 6 months ago

User session is not stored in a shared "memory" which could be used by all App Engine instances running the application.

upvoted 2 times

Ekramy_Elnaggar 1 year, 9 months ago

Selected Answer: A

Answer is A. The session variable is local to just a single instance.

Here's why:

* App Engine Instances: Google App Engine automatically scales your application by creating multiple instances to handle incoming traffic. Each instance runs independently and has its own memory space.

* Session Variable: In the provided code, the `sessions` variable is a simple Python dictionary stored in the memory of each instance. This means each instance has its own copy of the `sessions` data.

* The Problem: When a user logs in, their viewed articles are stored in the `sessions` variable of the instance that handled their request. If subsequent requests from the same user are routed to a different instance, that instance won't have any record of the previously viewed articles, causing them to be displayed again.

upvoted 6 times

Ekramy_Elnaggar 1 year, 9 months ago

Answer is A. The session variable is local to just a single instance. Here's why: * App Engine Instances: Google App Engine automatically scales your application by creating multiple instances to handle incoming traffic. Each instance runs independently and has its own memory space. * Session Variable: In the provided code, the `sessions` variable is a simple Python dictionary stored in the memory of each instance. This means each instance has its own copy of the `sessions` data. * The Problem: When a user logs in, their viewed articles are stored in the `sessions` variable of the instance that handled their request. If subsequent requests from the same user are routed to a different instance, that instance won't have any record of the previously viewed articles, causing them to be displayed again.how would you change the code?

upvoted 6 times

desertlotus1211 1 year, 8 months ago

how would you change the code?

upvoted 1 times

omermahgoub 1 year, 11 months ago

A The most likely cause of the issue described in the code is that the session variable is local to just a single instance. In this code, the session variable is defined as a local dictionary within the Flask application. This means that it is not shared across different instances of the application and

will not be persisted between requests. As a result, when the application is running on multiple instances, each instance will have its own local copy of the session variable, and users may see news articles that they have already viewed on other instances.

upvoted 4 times

omermahgoub 1 year, 11 months ago

A The most likely cause of the issue described in the code is that the session variable is local to just a single instance. In this code, the session variable is defined as a local dictionary within the Flask application. This means that it is not shared across different instances of the application and will not be persisted between requests. As a result, when the application is running on multiple instances, each instance will have its own local copy of the session variable, and users may see news articles that they have already viewed on other instances.To fix this issue, you could consider using a persistent storage solution, such as Cloud Datastore or Cloud SQL, to store the session data in a way that is shared across all instances of the application. This would allow you to maintain a consistent view of the session data for each user across all instances of the application. Other potential causes for this issue, such as modifying the URL of the API to prevent caching or setting the HTTP Expires header to -1 to stop caching, are not related to the issue described in the code and would not likely address the problem.

upvoted 4 times

omermahgoub 3 years, 8 months ago

To fix this issue, you could consider using a persistent storage solution, such as Cloud Datastore or Cloud SQL, to store the session data in a way that is shared across all instances of the application. This would allow you to maintain a consistent view of the session data for each user across all instances of the application. Other potential causes for this issue, such as modifying the URL of the API to prevent caching or setting the HTTP Expires header to -1 to stop caching, are not related to the issue described in the code and would not likely address the problem.

upvoted 2 times

Badri9898 1 year, 11 months ago

The most likely cause of the reported issue is that the session variable is local to just a single instance. In the code provided, the sessions variable is a dictionary that stores the viewed news articles for each user. However, this variable is only stored in memory on the instance that handles the request, and it is not shared between instances. Therefore, when a new request is handled by a different instance, it will not have access to the same session data, and the user may see previously viewed news articles. To solve this problem, a shared session management system should be used that can be accessed by all instances. Google App Engine provides a few options for session management, such as using Memcache or Cloud Datastore to store the session data. By using a shared session management system, all instances can access the same session data, and users will not see previously viewed news articles.

upvoted 8 times

kurili 1 year, 11 months ago

A.The sessions dictionary is used to store user-specific data, such as which news articles have been viewed. This dictionary is created as an in-memory variable within the Flask app. In a cloud environment, like Google App Engine, the application may be running on multiple instances, especially during peak loads. Since the sessions dictionary is stored in memory, it is local to each instance. This means that if a user is routed to a different instance (due to load balancing), their session data will not be available on that new instance, causing the application to serve news articles they've already seen.

upvoted 1 times

simonab23 3 years, 6 months ago

A is the right answer

upvoted 1 times

jay9114 3 years, 7 months ago

Where in the code does it show that the session variable is local to just a single instance?

upvoted 3 times

Amrit123_ 3 years, 8 months ago

A is correct

upvoted 1 times

angelumesh 3 years, 8 months ago

stateful variable should be in firestore (redis).

upvoted 1 times

Racinely 3 years, 9 months ago

I agree with ackdbd

upvoted 1 times

Mahmoud_E 3 years, 10 months ago

A is the correct answer app becoming stateful and it should not be sin case of app engine, cloud run and functions

upvoted 1 times

minmin2020 3 years, 10 months ago

A. The session variable is local to just a single instance The others are not relevant

upvoted 1 times

Question #5 *Topic 1*

An application development team believes their current logging tool will not meet their needs for their new cloud-based product. They want a better tool to capture errors and help them analyze their historical log data. You want to help them find a solution that meets their needs. What should you do?

A. Direct them to download and install the Google StackDriver logging agent

B. Send them a list of online resources about logging best practices

C. Help them define their requirements and assess viable logging tools

D. Help them upgrade their current tool to take advantage of any new features

Suggested Answer: *C*

Official Answer: C

Community vote distribution

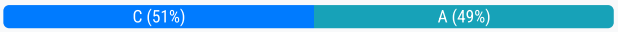

dummyemailforexam Highly Voted 6 years, 3 months ago

A. This is GCP exam. They will always promote their services. Not a third party solution.Couldn't DISagree more. There are multiple cloud-based solutions that involve parties in and out GCP, and for an Architect it's a very risky (and almost irresponsible) decision to simply ask them to install that if we don't know what they need. A solutions architect, in and out of Google, will ALWAYS have to consider requirements before proposing any solution. We don't even know why they think their current logging tool will not meet their needs, and while it seems "obvious" that Google will promote their own service, designing a solution that doesn't work will hurt Google's reputation in the long term. C is the right answer to me.I could not disagree with this. 100%AgreedRequirements before solutions.

upvoted 130 times

dummyemailforexam Highly Voted 6 years, 3 months ago

A. This is GCP exam. They will always promote their services. Not a third party solution.

upvoted 130 times

jrisl1991 2 years, 11 months ago

Couldn't DISagree more. There are multiple cloud-based solutions that involve parties in and out GCP, and for an Architect it's a very risky (and almost irresponsible) decision to simply ask them to install that if we don't know what they need. A solutions architect, in and out of Google, will ALWAYS have to consider requirements before proposing any solution. We don't even know why they think their current logging tool will not meet their needs, and while it seems "obvious" that Google will promote their own service, designing a solution that doesn't work will hurt Google's reputation in the long term. C is the right answer to me.

upvoted 13 times

zr79 3 years, 10 months ago

I could not disagree with this. 100%

upvoted 5 times

AzureDP900 3 years, 10 months ago

Agreed

upvoted 1 times

Moophoop 3 years, 4 months ago

Requirements before solutions.

upvoted 8 times

MeasService 6 years, 10 months ago

C should be the correct answer hereC is okWhat all viable logging you would suggest in this scenario! Cloud operations suite has everything.. cloud logging helps with many thing..In my mind this question is not meant for being a perfectionist but what would mostly work while option C is an approach that we will take questions ask about tool. Hence selecting A make sense to me. Thoughts!I would love to choose B, but need to keep my job..I'm surprised, @MeasService. I guess you had created the question an sugested ans. A. Then you wrote "C should be the correct answer here". Do you change your mind ?The Stackderiver agent currently known as Ops agent is the primary agent for collecting telemetry from your Compute Engine instances. It needs to be installed on GCP services such as GCE instances in order to collect logs from those instances and send them to cloud logging and monitoring. https://cloud.google.com/stackdriver/docs/solutions/agents/ops-agent/installation.Review the logging requirements and use existing logging utility. is not right. You know the requirements as you have found the current logging tool will not meet the needs for the new

cloud-based product. Install Google Cloud logging agent on all VMs. is the right answer. The Logging agent streams logs from common third-party applications and system software to Logging. It is a best practice to run the Logging agent on all your VM instances. The agent runs under both Linux and Windows. The Logging agent is compatible with both GCP Compute Engine instances as well as AWS EC2 instances. Google, through its partners, provides logging services for on-premise and hybrid cloud platforms consistently and predictably. Ref: https://cloud.google.com/logging/docs/agentWhy do you think they are using VM on GCE?This is wrong. This certification is meant to prepare you to be a cloud architect (focusing on GCP). This does not mean that you are going to recommend everything that has google on its name. You need to understa

upvoted 65 times

MeasService 6 years, 10 months ago

C should be the correct answer here

upvoted 65 times

tartar 6 years ago

C is ok

upvoted 10 times

tartar 6 years ago

C is okWhat all viable logging you would suggest in this scenario! Cloud operations suite has everything.. cloud logging helps with many thing..In my mind this question is not meant for being a perfectionist but what would mostly work while option C is an approach that we will take questions ask about tool. Hence selecting A make sense to me. Thoughts!I would love to choose B, but need to keep my job..

upvoted 10 times

Vika 5 years, 5 months ago

What all viable logging you would suggest in this scenario! Cloud operations suite has everything.. cloud logging helps with many thing..In my mind this question is not meant for being a perfectionist but what would mostly work while option C is an approach that we will take questions ask about tool. Hence selecting A make sense to me. Thoughts!

upvoted 2 times

tartar 6 years ago

I would love to choose B, but need to keep my job..

upvoted 16 times

lynx256 5 years, 4 months ago

I'm surprised, @MeasService. I guess you had created the question an sugested ans. A. Then you wrote "C should be the correct answer here". Do you change your mind ?

upvoted 7 times

aatt1122 3 years, 7 months ago

The Stackderiver agent currently known as Ops agent is the primary agent for collecting telemetry from your Compute Engine instances. It needs to be installed on GCP services such as GCE instances in order to collect logs from those instances and send them to cloud logging and monitoring. https://cloud.google.com/stackdriver/docs/solutions/agents/ops-agent/installation.

upvoted 1 times

techalik 1 year, 11 months ago

Review the logging requirements and use existing logging utility. is not right.

You know the requirements as you have found the current logging tool will not meet the needs for the new cloud-based product.

Install Google Cloud logging agent on all VMs. is the right answer.

The Logging agent streams logs from common third-party applications and system software to Logging. It is a best practice to run the Logging agent on all your VM instances. The agent runs under both Linux and Windows. The Logging agent is compatible with both GCP Compute Engine instances as well as AWS EC2 instances. Google, through its partners, provides logging services for on-premise and hybrid cloud platforms consistently and predictably.

Ref: https://cloud.google.com/logging/docs/agent

upvoted 4 times

techalik 1 year, 11 months ago

Review the logging requirements and use existing logging utility. is not right. You know the requirements as you have found the current logging tool will not meet the needs for the new cloud-based product. Install Google Cloud logging agent on all VMs. is the right answer. The Logging agent streams logs from common third-party applications and system software to Logging. It is a best practice to run the Logging agent on all your VM instances. The agent runs under both Linux and Windows. The Logging agent is compatible with both GCP Compute Engine instances as well as AWS

EC2 instances. Google, through its partners, provides logging services for on-premise and hybrid cloud platforms consistently and predictably. Ref: https://cloud.google.com/logging/docs/agentWhy do you think they are using VM on GCE?This is wrong. This certification is meant to prepare you to be a cloud architect (focusing on GCP). This does not mean that you are going to recommend everything that has google on its name. You need to understand client requirements first.Actually it does. A lot of questions guide you to the, in best case, fully managed GCP services. They should rename the certificate to "Professional Google Cloud Sales Person".

upvoted 4 times

okixavi 5 years, 8 months ago

Why do you think they are using VM on GCE?

upvoted 1 times

penelop 4 years, 11 months ago

This is wrong. This certification is meant to prepare you to be a cloud architect (focusing on GCP). This does not mean that you are going to recommend everything that has google on its name. You need to understand client requirements first.

upvoted 5 times

penelop 4 years, 11 months ago

This is wrong. This certification is meant to prepare you to be a cloud architect (focusing on GCP). This does not mean that you are going to recommend everything that has google on its name. You need to understand client requirements first.Actually it does. A lot of questions guide you to the, in best case, fully managed GCP services. They should rename the certificate to "Professional Google Cloud Sales Person".

upvoted 5 times

aceton999 4 years, 6 months ago

Actually it does. A lot of questions guide you to the, in best case, fully managed GCP services. They should rename the certificate to "Professional Google Cloud Sales Person".

upvoted 7 times

Wayne90040 Most Recent 3 months, 3 weeks ago

C is the answer.

upvoted 1 times

Johnb9682 4 months, 1 week ago

This is GCP exam. They will always promote their services. Not a third party solution.

upvoted 1 times

ch_rehan 7 months ago

Option C is the best practice because it ensures the solution is tailored to the team's specific technical and business requirements (e.g., log volume, retention periods, query performance, alerting, cost constraints, multi-cloud considerations, and integration with monitoring or analytics tools). Only after requirements are clearly defined can viable options be properly evaluated and compared.

upvoted 2 times

MissK1371 8 months, 3 weeks ago

google do never force their product. Def answer C

upvoted 1 times

mgrmek 1 year ago

a should be

upvoted 1 times

AssafKatz 1 year, 1 month ago

Although the managed resource is good, there is no indication that it satisfies the requirements. Google Exam may handle it differently, I agree.

upvoted 1 times

shivaky 1 year, 2 months ago

Correct Answer: C. Help them define their requirements and assess viable logging tools. Explanation: - Before selecting a logging tool, it's critical to understand the team's specific needs, including error tracking, historical log analysis, scalability, and integration with their cloud-based product. - Assessing viable tools ensures they choose a solution that meets their technical and business requirements rather than just installing a predefined option. - This approach avoids unnecessary upgrades and ensures they adopt a tool that fits their workflow and future growth.

upvoted 2 times

rajrocks171 1 year, 2 months ago

I would go with C. Understanding needs crucial here before suggesting a tool or approach

upvoted 1 times

rinkujacob 1 year, 2 months ago

As an architect requirements needs to be clear before proposing solution. Cloud Logging and Error Reporting could be possible options.

upvoted 2 times

gtc009 1 year, 3 months ago

Need to evaluate their requirements befor choosing a Tool

upvoted 1 times

mahi_h 1 year, 3 months ago

Answer is A, it's always google products and services first in this exam.

upvoted 1 times

Dagged 1 year, 4 months ago

C makemore sense, but as it is GCP exam after all, they wil promote thier own tools .A would be correct

upvoted 1 times

Sib09 1 year, 5 months ago

We need to define requirements

upvoted 1 times

PetarMarinkovic 1 year, 5 months ago

Come on guys, this is GCP

upvoted 2 times

cloud_rider 1 year, 5 months ago

A is the right answer.

upvoted 1 times

Question #6 *Topic 1*

You need to reduce the number of unplanned rollbacks of erroneous production deployments in your company's web hosting platform.
Improvement to the QA/
Test processes accomplished an 80% reduction.
Which additional two approaches can you take to further reduce the rollbacks? (Choose two.)

A. Introduce a green-blue deployment model

B. Replace the QA environment with canary releases

C. Fragment the monolithic platform into microservices

D. Reduce the platform's dependency on relational database systems

E. Replace the platform's relational database systems with a NoSQL database

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

Community vote distribution

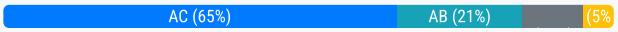

JustJack21 Highly Voted 4 years, 11 months ago
D) and E) are pointless in this context. C) is certainly a good practice. Now between A) and B) A) Blue green deployment is an application release model that gradually transfers user traffic from a previous version of an app or microservice to a nearly identical new release—both of which are running in production. c) In software, a canary process is usually the first instance that receives live production traffic about a new configuration update, either a binary or configuration rollout. The new release only goes to the canary at first. The fact that the canary handles real user traffic is key: if it breaks, real users get affected, so canarying should be the first step in your deployment process, as opposed to the last step in testing in production. " While both green-blue and canary releases are useful, B) suggests "replacing QA" with canary releases - which is not good. QA got the issue down by 80%. Hence A) and C)
upvoted 71 times

jdpinto Highly Voted 5 years, 2 months ago
A & C for me
upvoted 35 times

hitmax87 Most Recent 2 months, 3 weeks ago
Who would choose the C option?))) Redesigning solution architecture requires migration plan and at least 6 months of development which definitely will cause new bugs. No one will go this way until delivery process is really sluggish and scalability is poor. There should be solid reasons, not just reducing 20% of rollbacks)) Come on, guys.
upvoted 1 times

Johnb9682 4 months, 1 week ago
A, C Seems correct
upvoted 1 times

ch_rehan 6 months, 3 weeks ago
Even the question is not a good question in context of it's cohesion but the correct options are A & C. Blue green deployments allow instant rollback by switching traffic, which directly reduces production impact. Breaking a monolith into microservices limits blast radius so failures affect smaller components instead of the whole platform.
upvoted 1 times

sky1024 8 months ago
Canary is for beta testing before rolling out a feature to end users in production not for lower envs. are these questions still valid and for renewal exam?
upvoted 1 times

29d063d 8 months, 1 week ago
A. Introduce a green-blue deployment model: This allows for an instantaneous, non-disruptive rollback by flipping traffic back to the known good "green" environment if the newly deployed "blue" environment fails. B. Replace the QA environment with canary releases: This minimizes risk by deploying new code to a small percentage of real users and infrastructure, catching production-specific errors before a full rollout
upvoted 2 times

41168b9 8 months, 2 weeks ago

A.) This is an excellent strategy because it eliminates downtime during deployment and provides an instant, low-risk rollback mechanism. B.) While not replacing the entire QA process, implementing Canary Releases as the final stage of deployment is highly effective. It allows you to test the new code on a small subset of actual users in the real production environment.

upvoted 1 times

Vzl 11 months, 1 week ago

We "reduce the number of unplanned rollbacks of erroneous production deployments". Both blue-green and canary deployment help reduce the"unplanned" "erroneous" production deployment

upvoted 2 times

mgrmek 1 year ago

a and b sould be

upvoted 1 times

Klastorm 1 year, 2 months ago

A) Enables immediate rollback by switching traffic back to the blue environment if issues arise B) Limits the blast radius of issues, enabling rapid fixes without full rollbacks I can't understand why many of you voted C.. C: Fragmentation improves scalability but doesn't directly address rollback mechanisms. It's for sure a good practice but it's not related to the subject of question. This is my point of view.You either use GREEN-BLUE or CANARY. The Q&A has already reduced rollbacks by 80%. Question states what else can we do. Undoing what's working is not ideal...

upvoted 1 times

Klastorm 1 year, 2 months ago

Selected Answer: AB

A) Enables immediate rollback by switching traffic back to the blue environment if issues arise B) Limits the blast radius of issues, enabling rapid fixes without full rollbacks

I can't understand why many of you voted C..

C: Fragmentation improves scalability but doesn't directly address rollback mechanisms. It's for sure a good practice but it's not related to the subject of question. This is my point of view.

upvoted 1 times

ast3citos 9 months, 3 weeks ago

You either use GREEN-BLUE or CANARY. The Q&A has already reduced rollbacks by 80%. Question states what else can we do. Undoing what's working is not ideal...

upvoted 2 times

mahi_h 1 year, 3 months ago

A&C are suitable approaches

upvoted 1 times

befeyac438 1 year, 3 months ago

You can run blue green, and a canary inside green to move the rollback towards zero.

upvoted 1 times

izekc 1 year, 5 months ago

A is not correct

upvoted 1 times

halifax 1 year, 6 months ago

Fragmenting the monolithic platform into microservices (option C) would take a few months to a few years - - This is a long-term solution.

upvoted 3 times

hpf97 1 year, 6 months ago

Selected Answer: AC

D&E : Not related to the problem A - B/G deployment does not improve quality. In B/G, your deploy the new application in parallel of the old one, and fully move to it when ready. So if there is an error, you will have to go back to the old one. So not really a rollback which would mean downgrade the system. B - With Canary, you will gradually move to the new one if everything is OK, or stay on the old one. So no rollback either..... but as it is say to replace QA environment, I will not keep this answer C : Best practice to reduce change scope and change size. So A & C

upvoted 1 times

hpf97 1 year, 6 months ago

D&E : Not related to the problem A - B/G deployment does not improve quality. In B/G, your deploy the new application in parallel of the old one, and fully move to it when ready. So if there is an error, you will have to go back to the old one. So not really a rollback which would mean downgrade the system. B - With Canary, you will gradually move to the new one if everything is OK, or stay on the old one. So no rollback either..... but as it is say to

replace QA environment, I will not keep this answer C : Best practice to reduce change scope and change size. So A & CC would mean increase in number of deployments (because each service has its own deployment) and hence implied increase in number of rollbacks , no? With monolith, you just had the 1 deployment/rollback

upvoted 1 times

pm2da 1 year, 6 months ago

C would mean increase in number of deployments (because each service has its own deployment) and hence implied increase in number of rollbacks , no? With monolith, you just had the 1 deployment/rollback

upvoted 2 times

alihabib 1 year, 6 months ago

D & E is unrelated, C is controversial, it requires a rewrite of application code completely. I would go for A & B, as it more focused on Release Management & provides continuous feedback

upvoted 1 times

Question #7 *Topic 1*

To reduce costs, the Director of Engineering has required all developers to move their development infrastructure resources from on-premises virtual machines
(VMs) to Google Cloud Platform. These resources go through multiple start/stop events during the day and require state to persist. You have been asked to design the process of running a development environment in Google Cloud while providing cost visibility to the finance department.
Which two steps should you take? (Choose two.)

A. Use the - -no-auto-delete flag on all persistent disks and stop the VM

B. Use the - -auto-delete flag on all persistent disks and terminate the VM

C. Apply VM CPU utilization label and include it in the BigQuery billing export

D. Use Google BigQuery billing export and labels to associate cost to groups

E. Store all state into local SSD, snapshot the persistent disks, and terminate the VM

F. Store all state in Google Cloud Storage, snapshot the persistent disks, and terminate the VM

Suggested Answer: *A, D*
Official Answer: AD

Community vote distribution

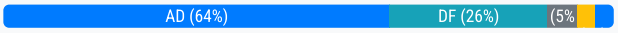

rishab86 Highly Voted 5 years, 2 months ago
A and D looks correct as per https://cloud.google.com/sdk/gcloud/reference/compute/instances/set-disk-auto-delete#--auto-delete ; https://cloud.google.com/billing/docs/how-to/export-data-bigquery-no-auto-delete flag does not have effect on the state of the application. I believe D and F are correct ANS, https://cloud.google.com/compute/docs/instances/stop-start-instance
upvoted 24 times

rishab86 Highly Voted 5 years, 2 months ago
A and D looks correct as per https://cloud.google.com/sdk/gcloud/reference/compute/instances/set-disk-auto-delete#--auto-delete ; https://cloud.google.com/billing/docs/how-to/export-data-bigquery
upvoted 24 times

RKS_2021 5 years, 1 month ago
-no-auto-delete flag does not have effect on the state of the application. I believe D and F are correct ANS, https://cloud.google.com/compute/docs/instances/stop-start-instance
upvoted 4 times

kopper2019 5 years, 1 month ago
A and D, are OK
upvoted 6 times

Upender_PDE Most Recent 8 months, 1 week ago
A -- > ensures state persists through stop/start cycles using persistent disks D -- > provides the cost visibility finance needs through BigQuery exports and resource labels
upvoted 1 times

hpf97 1 year, 6 months ago
A / B : the option only apply when VM are deleted (terminated ?), or here we are talking about stop/start See https://cloud.google.com/compute/docs/samples/compute-disk-autodelete-change C : Do even know if it is possible D : In order to bring "visibility on the cost" E : local disks does not persist F : complex solution, but it works So D & F
upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago
Answer is : A& D A. Use the --no-auto-delete flag on all persistent disks and stop the VM: 1. Cost Savings: When you stop a VM, you only pay for the persistent disks attached to it. The --no-auto-delete flag ensures that the disks remain available even when the VM is stopped, preserving the developers' work and avoiding the cost of recreating the environment from scratch each time. 2. State Persistence: This approach ensures that the development environment's state is saved on the persistent disk, allowing developers to resume their work seamlessly when they restart the VM. D. Use Google BigQuery billing export and labels to associate cost to groups: 1. Cost Visibility: BigQuery billing export allows you to analyze your Google

Cloud costs in detail. By applying labels to your resources (e.g., "environment: development", "team: frontend"), you can categorize and track costs associated with different development groups.

upvoted 3 times

Hungdv 2 years ago

Choose A and D

upvoted 1 times

Sephethus 2 years, 2 months ago

Another confusing question because I took the "these machines go through multiple stop/starts during the day" as a part of the migration, not as a part of daily functionality, so none of the answers other than D made much sense to me. People need to word the questions better on tests and give more than enough context or people like me are going to get confused, second guess our answers, and fail.

upvoted 1 times

afsarkhan 2 years, 3 months ago

F is too complex solution to solve this problem. E local SSD does not persist on termination of vm so this is also a wrong option A, B suggest persistent disk but I think A makes better sense. So my answer is A and D

upvoted 1 times

44eacc1 2 years, 3 months ago

E wrong: Scenarios where Compute Engine does not persist Local SSD data Data on Local SSD disks does not persist through the following events: If you shut down the guest operating system and force the VM to stop.

upvoted 1 times

shikha344 2 years, 5 months ago

Hi all, i am trying to view questions from 135.But i cannot access the page as it is asking for contributer access.Is it same for everyone?

upvoted 1 times

d0094d6 2 years, 6 months ago

From the GCP Practice exam... A and D

upvoted 2 times

Teckexam 2 years, 7 months ago

Based on documentation A is correct https://cloud.google.com/sdk/gcloud/reference/compute/instances/set-disk-auto-delete Also the documentation clearly states that this flag will be help retain the state when VM is started/stopped. https://cloud.google.com/blog/products/storage-data-transfer/save-money-by-stopping-and-starting-compute-engine-instances-on-schedule For cost visibility option D is correct.

upvoted 1 times

hzaoui 2 years, 7 months ago

D. Use Google BigQuery billing export and labels to associate cost to groups: This offers granular cost visibility across various development teams or projects through BigQuery data analysis. Combining labels with billing export allows you to associate resource consumption with specific groups, enabling chargeback mechanisms and fostering cost accountability. E. Store all state into local SSD, snapshot the persistent disks, and terminate the VM: This option minimizes ongoing cost by utilizing low-cost local SSD for active state during runtime and terminating the VM when development isn't ongoing. Snapshots offer quick restoration back to the latest state without incurring persistent disk charges during downtime.

upvoted 1 times

kshlgpt 2 years, 7 months ago

DF. This is the question in google practice test.

upvoted 2 times

cfigueiredo 2 years, 7 months ago

D & F for me

upvoted 2 times

Prakzz 2 years, 10 months ago

How can D ever be right coz it's Bigquery Billing Export and question is about VM billing

upvoted 1 times

Prakzz 2 years, 10 months ago

How can D ever be right coz it's Bigquery Billing Export and question is about VM billingThis means you can export the billing to the BQ and then do analysis.

upvoted 1 times

ArtistS 2 years, 9 months ago

This means you can export the billing to the BQ and then do analysis.

upvoted 1 times

AdityaGupta 2 years, 10 months ago

A is correct, Use of persistent disk mean the data is preserved even after restart. -np-auto-delete on persistent disk means pesistent disk won't be deleted when VM is deleted. D is correct, becuase second part of question asks for billing report to finanace department and Label and BQ helps in cost analysis.

upvoted 4 times

Question #8 *Topic 1*

Your company wants to track whether someone is present in a meeting room reserved for a scheduled meeting. There are 1000 meeting rooms across 5 offices on 3 continents. Each room is equipped with a motion sensor that reports its status every second. The data from the motion detector includes only a sensor ID and several different discrete items of information. Analysts will use this data, together with information about account owners and office locations.
Which database type should you use?

A. Flat file

B. NoSQL

C. Relational

D. Blobstore

Suggested Answer: *B*
Relational databases were not designed to cope with the scale and agility challenges that face modern applications, nor were they built to take advantage of the commodity storage and processing power available today. NoSQL fits well for: ✑ Developers are working with applications that create massive volumes of new, rapidly changing data types ג€" structured, semi-structured, unstructured and polymorphic data. Incorrect Answers: D: The Blobstore API allows your application to serve data objects, called blobs, that are much larger than the size allowed for objects in the Datastore service. Blobs are useful for serving large files, such as video or image files, and for allowing users to upload large data files. Reference: https://www.mongodb.com/nosql-explained

Community vote distribution
B (83%) C (17%)

clouddude Highly Voted 6 years, 3 months ago
I'll go with B. This is time series data. We also have no idea what kinds of data are being captured so it doesn't appear structurd. A does not seem reasonable because a flat file is not easy to query and analyze. B seems reasonable because this accommodates unstructured data. C seems unreasonable because we have no idea on the structure of the data. D seems unreasonable beacause there is no such Google database type.
upvoted 35 times

hitmax87 Most Recent 2 months, 3 weeks ago
SQL would work here. But NoSQL will be faster. Because: 1. We mostly append data, we don`t update or delete data. We also don`t have any relations. 2. Analysts typically request bulks of data, so that it will work faster with warehouse than with SQL. I dont think DB distribution is meaningful here, nowadays there are solutions that allow to replicate both SQL and NoSQL DBs across regions with either strong and eventual consistency.
upvoted 1 times

41168b9 8 months, 2 weeks ago
A Relational Database Management System (RDBMS) is the ideal choice because the application involves combining or joining several distinct, structured data sets to answer analytical questions. Data Structure and Relationships The system must handle and link at least three different types of structured information: Motion Sensor Data: High-volume time series data (Sensor ID, timestamp, presence status). Room/Office Data: Static data linking the Sensor ID to the Room Name, Office Location, and Continent. Meeting/Account Data: Static data linking Room Name and Time Slot to the Account Owner (who reserved the room).
upvoted 1 times

JoseCloudEng1994 11 months, 3 weeks ago
It is B, From the start flat file cannot be an answer Blobstore (I believe that is particular to AppEngine but its really old) so I guess they are referring to Cloud Storage which stores things in Blobs. So yeah, not this one either, no way to query it Then we know its semi-structured date so its got to be B
upvoted 1 times

teofil 1 year ago
While a NoSQL database would be excellent at ingesting the high volume of sensor data, it would make the required analytical queries that join different data entities significantly more complex and inefficient to perform. A modern, managed relational database (like Google Cloud SQL or Spanner) is designed to handle both the required transaction volume and the complex analytical queries.
upvoted 1 times

mahi_h 1 year, 2 months ago
B is the suitable option for unstructured data
upvoted 1 times

hpf97 1 year, 6 months ago

A / D : No very good in order to exploit the data C : Relational DB would have been a good candidate for managing sensor data only. But it is said to managing other data type which are not detailled. So B is the best

upvoted 1 times

ddatta 1 year, 8 months ago

We don't know the data type. Only nosql make sense.

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

Answer is B. 1. High Volume and Velocity of Data: You have 1000 rooms reporting data every second, resulting in a massive amount of data with high velocity. NoSQL databases are designed to handle this kind of volume and speed efficiently. 2. Simple Data Structure: The data from the motion sensor is relatively simple (sensor ID and discrete information). NoSQL databases are well-suited for storing and processing this type of data without the need for complex schemas. 3. Flexible Schema: NoSQL databases offer schema flexibility, allowing you to easily adapt to changes in the data structure if needed. This is important as your tracking requirements might evolve over time. 4.Scalability: NoSQL databases are highly scalable, making it easy to accommodate future growth in the number of meeting rooms or data volume.

upvoted 3 times

ashishdwi007 2 years, 7 months ago

With frequencies of data (per second), the best case would be using pub/sub and NoSQL. Relational DB/BlobStore/FlatFile are not good for Near realtime data.

upvoted 2 times

hzaoui 2 years, 7 months ago

C. Relational database. Here's why: Scalability: A relational database can handle the data volume from 1000 sensors reporting every second effectively. Structure: It provides a well-defined schema for organizing data like sensor ID, timestamp, motion status, account owner, and office location, making it easily queryable and understandable for analysts. Relationships: It allows establishing relationships between tables, such as linking sensor data to specific meeting rooms and their corresponding owners and locations. This facilitates analyses involving multiple data sources. Flexibility: Relational databases offer flexibility for expanding data collection beyond motion sensors in the future to include other sensor types or meeting room details.

upvoted 2 times

_kartik_raj 2 years, 10 months ago

B, It is

upvoted 1 times

AdityaGupta 2 years, 10 months ago

Unstructured realtime aata

upvoted 1 times

ChinaSailor 2 years, 11 months ago

b [You need seperate fields and keys -- you do not need to relate them

upvoted 1 times

heretolearnazure 2 years, 12 months ago

NOSQL DB's are meant for these kind of workloads

upvoted 1 times

BiddlyBdoyng 3 years, 1 month ago

The requirement to join the data to other data sets implies RDBMS. BigQuery can handle 1GB/s when streaming inserts, I doubt these 1000 sensors will send that much data. Bigtable seems over the top and not able to fulfil all the requirements.

upvoted 2 times

Deb2293 3 years, 5 months ago

This will be time-series data. The best DB would be a Big Table (also sensorID can be used in the row key for faster retrieval of data)

upvoted 3 times

Question #9 *Topic 1*

You set up an autoscaling instance group to serve web traffic for an upcoming launch. After configuring the instance group as a backend service to an HTTP(S) load balancer, you notice that virtual machine (VM) instances are being terminated and re-launched every minute. The instances do not have a public IP address.
You have verified the appropriate web response is coming from each instance using the curl command. You want to ensure the backend is configured correctly.
What should you do?

A. Ensure that a firewall rules exists to allow source traffic on HTTP/HTTPS to reach the load balancer.

B. Assign a public IP to each instance and configure a firewall rule to allow the load balancer to reach the instance public IP.

C. Ensure that a firewall rule exists to allow load balancer health checks to reach the instances in the instance group.

D. Create a tag on each instance with the name of the load balancer. Configure a firewall rule with the name of the load balancer as the source and the instance tag as the destination.

Suggested Answer: *C*
The best practice when configuration a health check is to check health and serve traffic on the same port. However, it is possible to perform health checks on one port, but serve traffic on another. If you do use two different ports, ensure that firewall rules and services running on instances are configured appropriately. If you run health checks and serve traffic on the same port, but decide to switch ports at some point, be sure to update both the backend service and the health check. Backend services that do not have a valid global forwarding rule referencing it will not be health checked and will have no health status. Reference: https://cloud.google.com/compute/docs/load-balancing/http/backend-service

Community vote distribution
C (100%)

Eroc Highly Voted 1 year, 11 months ago
"A" and "B" wouldn't turn the VMs on or off, it would jsut prevent traffic. "C" would turn them off if the health check is configured to terminate the VM is it fails. "D" is the start of a pseudo health check without any logic, so it also isn't an answer because it is like "A" and "B". Correct Answer: "C"C is okC because terminated and relaunch.... something wrong with HC.agreed and C is right
upvoted 38 times

Eroc Highly Voted 1 year, 11 months ago
"A" and "B" wouldn't turn the VMs on or off, it would jsut prevent traffic. "C" would turn them off if the health check is configured to terminate the VM is it fails. "D" is the start of a pseudo health check without any logic, so it also isn't an answer because it is like "A" and "B". Correct Answer: "C"
upvoted 38 times

tartar 6 years ago
C is ok
upvoted 14 times

nitinz 5 years, 5 months ago
C because terminated and relaunch.... something wrong with HC.
upvoted 6 times

AzureDP900 3 years, 10 months ago
agreed and C is right
upvoted 2 times

TheCloudBoy77 4 years, 9 months ago
A. Ensure that a firewall rules exists to allow source traffic on HTTP/HTTPS to reach the load balancer. >> not correct, load balancer is not the issue here. B. Assign a public IP to each instance and configure a firewall rule to allow the load balancer to reach the instance public IP. >> defeats the purpose of getting load balancers , not correct C. Ensure that a firewall rule exists to allow load balancer health checks to reach the instances in the instance group.>> Correct. if using different port then appropriate FW rule need to be setup to ensure LB can reach backend instances for healthcheck. if healthcheck traffic is blcked, instances will be marked unhealthy and will be restarted. D. Create a tag on each instance with the name of the load balancer. Configure a firewall rule with the name of the load balancer as the source and the instance tag as the destination.>> tagging is not useful here as the instance is not the source of traffic, just the port need to be opened on FW.
upvoted 7 times

Upender_PDE Most Recent 8 months, 1 week ago

Why C is correct Google HTTP(S) Load Balancers perform health checks from specific IP ranges You must allow traffic from these IP ranges to your backend VMs: 35.191.0.0/16 130.211.0.0/22 If not allowed: Health checks fail Instances are continuously recreated Creating a firewall rule to allow health check traffic fixes this.

upvoted 2 times

hpf97 1 year, 6 months ago

The only answer which deals with the problem of the frequent stop / start of instances in a MIG.

upvoted 1 times

Houssemonline 1 year, 7 months ago

C is correct because health check failures lead to a VM being marked unhealthy and can result in termination if the health check continues to fail. Because you have already verified that the instances are functioning properly, the next step would be to determine why the health check is continuously failing.

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Health Checks are Essential: Load balancers rely on health checks to determine if instances in the backend pool are healthy and able to serve traffic. If the health checks fail, the load balancer assumes the instance is unhealthy and terminates it, leading to the constant cycling you're observing. 2. Firewall Rules for Health Checks: Even though your instances don't have public IPs, the load balancer needs to communicate with them through internal IPs for health checks. Firewall rules must be configured to allow this communication. 3. How Health Checks Work: The load balancer sends requests (e.g., HTTP, HTTPS, TCP) to the instances on a specific port and expects a certain response. The firewall needs to allow these requests to reach the instances and the responses to return to the load balancer.

upvoted 2 times

subramanyam46 2 years, 5 months ago

c is right

upvoted 1 times

hzaoui 2 years, 7 months ago

C is correct

upvoted 1 times

yas_cloud 2 years, 7 months ago

Most likely the problem of instances terminating is with the threshold settings on the health check. It's thinking too sooner that some VMs can be terminated due to less load.

upvoted 1 times

Arun_m_123 2 years, 10 months ago

One thing that i couldn't understand is - How VMs getting terminated and relaunched for not setting health-checks in the load balancer ? how will that affect VM's uptime ?

upvoted 2 times

AdityaGupta 2 years, 10 months ago

If the healthcheck is not successful, it will keep on re-creating the instances in MIG.

upvoted 1 times

angelumesh 3 years, 8 months ago

C (LB Health check should be taken care of)

upvoted 2 times

Mahmoud_E 3 years, 10 months ago

C is the correct answer

upvoted 2 times

minmin2020 3 years, 10 months ago

C. Ensure that a firewall rule exists to allow load balancer health checks to reach the instances in the instance group.

upvoted 3 times

YAS007 4 years, 3 months ago

answer C: https://cloud.google.com/load-balancing/docs/health-check-concepts#ip-ranges

upvoted 2 times

AWS56 4 years, 6 months ago

C is corect.

upvoted 2 times

OrangeTiger 4 years, 7 months ago

C is corect.

upvoted 1 times

OrangeTiger 4 years, 7 months ago

C is corect.' (VM) instances are being terminated and re-launched every minute. ' Isn't it because the health check is failing. A & D Maybe aleady there.curl command passed. B What are you doing. Absolutely no.

upvoted 1 times

OrangeTiger 4 years, 7 months ago

' (VM) instances are being terminated and re-launched every minute. ' Isn't it because the health check is failing. A & D Maybe aleady there.curl command passed. B What are you doing. Absolutely no.

upvoted 2 times

Question #10 *Topic 1*

You write a Python script to connect to Google BigQuery from a Google Compute Engine virtual machine. The script is printing errors that it cannot connect to
BigQuery.
What should you do to fix the script?

A. Install the latest BigQuery API client library for Python

B. Run your script on a new virtual machine with the BigQuery access scope enabled

C. Create a new service account with BigQuery access and execute your script with that user

D. Install the bq component for gcloud with the command gcloud components install bq.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
C (73%) B (24%) (3

kalschi Highly Voted 1 year, 11 months ago
A - If client library was not installed, the python scripts won't run - since the question states the script reports "cannot connect" - the client library must have been installed. so it's B or C.

B - https://cloud.google.com/bigquery/docs/authorization an access scope is how your client application retrieve access_token with access permission in OAuth when you want to access services via API call - in this case, it is possible that the python script use an API call instead of library, if this is true, then access scope is required. client library requires no access scope (as it does not go through OAuth)

C - service account is Google Cloud's best practice So prefer C.
upvoted 102 times

kalschi Highly Voted 1 year, 11 months ago
A - If client library was not installed, the python scripts won't run - since the question states the script reports "cannot connect" - the client library must have been installed. so it's B or C. B - https://cloud.google.com/bigquery/docs/authorization an access scope is how your client application retrieve access_token with access permission in OAuth when you want to access services via API call - in this case, it is possible that the python script use an API call instead of library, if this is true, then access scope is required. client library requires no access scope (as it does not go through OAuth) C - service account is Google Cloud's best practice So prefer C.agreed to comment here . C seems like a good optionagree access scope is enabled by default https://cloud.google.com/bigquery/docs/authorization#authenticate_with_oauth_20 If you use the BigQuery client libraries, you do not need this information, as this is done for you automatically.Might be an old versionRight.. B might be right only if C was not the option. so i will go with C as it's recommended google's best practice.
upvoted 102 times

Vika 5 years, 5 months ago
agreed to comment here . C seems like a good option
upvoted 4 times

MQQNB 3 years, 12 months ago
agree access scope is enabled by default https://cloud.google.com/bigquery/docs/authorization#authenticate_with_oauth_20 If you use the BigQuery client libraries, you do not need this information, as this is done for you automatically.
upvoted 2 times

Musk 6 years, 2 months ago
Might be an old version
upvoted 4 times

gauravagrawal 5 years, 6 months ago
Right.. B might be right only if C was not the option. so i will go with C as it's recommended google's best practice.
upvoted 4 times

KouShikyou 6 years, 10 months ago
Why not B? It looks better for me.

upvoted 14 times

KouShikyou 6 years, 10 months ago

Why not B? It looks better for me.C is okSorry, B is ok. You can create service account, add user to service account, and grant the user role as Service Account User. You still need to enable BigQuery scope to make the Python script running the instance to access BigQuery.Stop confusing people, B) doesn't make any sense. Why would you use or create a whole new VM just because of a permission issue? If anything, just stop the instance and edit the scope of the default Compute Service Account and grant it the role through IAM. C) is the most appropriate answer since you can only set scopes of the default Compute Service Account, if you're using any other, there's no scope option - its access is dictated strictly by IAM in such scenario. So C) is the answer: Stop the VM, change the Service Account with the appropriate permissions and done. B) would still need to have permission the set through IAM & Admin, the scope isn't enough with the default Compute Service Account.cloud guy1, relax. tartar is the hero for google cloud and if you read his answer, he explains the service account user's role granting on this one as that is the best practiceC, no brainer. You need SA for using API period. Thats where your start your troubleshooting.I stand corrected, B you need to have scope. It is union of Scope + Service Account. If scope is not there, you are screwed anyways.Configure the Python API to use a service account with relevant BigQuery access enabled. is the right answer. It is likely that this service account this script is running under does not have the permissions to connect to BigQuery and that could be causing issues. You can prevent these by using a service account that has the necessary roles to access BigQuery. Ref: https://cloud.google.com/bigquery/docs/reference/libraries#cloud-console A service account is a special kind of account used by an application or a virtual machine (VM) instance, not a person. Ref: https://cloud.google.com/iam/docs/service-accounts

upvoted 14 times

tartar 6 years ago

C is ok

upvoted 11 times

tartar 6 years ago

C is okSorry, B is ok. You can create service account, add user to service account, and grant the user role as Service Account User. You still need to enable BigQuery scope to make the Python script running the instance to access BigQuery.Stop confusing people, B) doesn't make any sense. Why would you use or create a whole new VM just because of a permission issue? If anything, just stop the instance and edit the scope of the default Compute Service Account and grant it the role through IAM. C) is the most appropriate answer since you can only set scopes of the default Compute Service Account, if you're using any other, there's no scope option - its access is dictated strictly by IAM in such scenario. So C) is the answer: Stop the VM, change the Service Account with the appropriate permissions and done. B) would still need to have permission the set through IAM & Admin, the scope isn't enough with the default Compute Service Account.cloud guy1, relax. tartar is the hero for google cloud and if you read his answer, he explains the service account user's role granting on this one as that is the best practice

upvoted 11 times

tartar 6 years ago

Sorry, B is ok. You can create service account, add user to service account, and grant the user role as Service Account User. You still need to enable BigQuery scope to make the Python script running the instance to access BigQuery.Stop confusing people, B) doesn't make any sense. Why would you use or create a whole new VM just because of a permission issue? If anything, just stop the instance and edit the scope of the default Compute Service Account and grant it the role through IAM. C) is the most appropriate answer since you can only set scopes of the default Compute Service Account, if you're using any other, there's no scope option - its access is dictated strictly by IAM in such scenario. So C) is the answer: Stop the VM, change the Service Account with the appropriate permissions and done. B) would still need to have permission the set through IAM & Admin, the scope isn't enough with the default Compute Service Account.cloud guy1, relax. tartar is the hero for google cloud and if you read his answer, he explains the service account user's role granting on this one as that is the best practice

upvoted 16 times

tartar 6 years ago

Sorry, B is ok. You can create service account, add user to service account, and grant the user role as Service Account User. You still need to enable BigQuery scope to make the Python script running the instance to access BigQuery.

upvoted 16 times

cloudguy1 6 years ago

Stop confusing people, B) doesn't make any sense. Why would you use or create a whole new VM just because of a permission issue? If anything, just stop the instance and edit the scope of the default Compute Service Account and grant it the role through IAM. C) is the most appropriate answer since you can only set scopes of the default Compute Service Account, if you're using any other, there's no scope option - its access is dictated strictly by IAM in such scenario. So C) is the answer: Stop the VM, change the Service Account with the appropriate permissions and done. B) would still need to have permission the set through IAM & Admin, the scope isn't enough with the default Compute Service Account.cloud guy1, relax. tartar is the hero for google cloud and if you read his answer, he explains the service account user's role granting on this one as that is the best practice

upvoted 37 times

cloudguy1 6 years ago

Stop confusing people, B) doesn't make any sense. Why would you use or create a whole new VM just because of a permission issue? If anything, just stop the instance and edit the scope of the default Compute Service Account and grant it the role through IAM. C) is the most appropriate answer

since you can only set scopes of the default Compute Service Account, if you're using any other, there's no scope option - its access is dictated strictly by IAM in such scenario. So C) is the answer: Stop the VM, change the Service Account with the appropriate permissions and done. B) would still need to have permission the set through IAM & Admin, the scope isn't enough with the default Compute Service Account.

upvoted 37 times

certificatores 5 years, 9 months ago

cloud guy1, relax. tartar is the hero for google cloud and if you read his answer, he explains the service account user's role granting on this one as that is the best practice

upvoted 4 times

nitinz 5 years, 5 months ago

C, no brainer. You need SA for using API period. Thats where your start your troubleshooting.

upvoted 6 times

nitinz 5 years, 5 months ago

C, no brainer. You need SA for using API period. Thats where your start your troubleshooting.I stand corrected, B you need to have scope. It is union of Scope + Service Account. If scope is not there, you are screwed anyways.

upvoted 6 times

nitinz 5 years, 5 months ago

I stand corrected, B you need to have scope. It is union of Scope + Service Account. If scope is not there, you are screwed anyways.

upvoted 2 times

techalik 5 years, 8 months ago

Configure the Python API to use a service account with relevant BigQuery access enabled. is the right answer. It is likely that this service account this script is running under does not have the permissions to connect to BigQuery and that could be causing issues. You can prevent these by using a service account that has the necessary roles to access BigQuery. Ref: https://cloud.google.com/bigquery/docs/reference/libraries#cloud-console A service account is a special kind of account used by an application or a virtual machine (VM) instance, not a person. Ref: https://cloud.google.com/iam/docs/service-accounts

upvoted 6 times

Ral17 Most Recent 4 months, 3 weeks ago

The error is a connection error from the VM itself, not an authentication error. This strongly signals the VM's access scope doesn't include the BigQuery API. Access scopes are set at VM creation time and cannot be changed on a running VM. Creating a new service account with BigQuery access fixes IAM level permissions if VM's access scope doesn't include BigQuery, the connection still fails.

upvoted 1 times

ahmed75.zahran00 6 months, 2 weeks ago

access scope enabled means the vm created with service account access. gcloud compute instances create INSTANCE_NAME --zone=ZONE --service-account=SERVICE_ACCOUNT_EMAIL https://docs.cloud.google.com/bigquery/docs/authentication

upvoted 1 times

ch_rehan 6 months, 3 weeks ago

B or C both are correct but B is legacy approach & C is the modern one. So go for the latest one.

upvoted 1 times

elon_musk 11 months, 4 weeks ago

Access Scopes are a legacy but still required mechanism that defines which APIs a VM is allowed to call. If a VM is created without the bigquery or the full cloud-platform scope, it is physically incapable of getting the necessary credentials to connect to the BigQuery API. Since access scopes cannot be changed on a running VM (you must stop the VM to edit them), and historically could not be changed at all, creating a new VM with the correct scopes is the most direct solution presented. The error "cannot connect" strongly suggests the issue is at this fundamental level, not a more specific "permission denied" error that would point to an IAM issue.

upvoted 2 times

YC_Lorenzo 1 year, 3 months ago

C is not the best option as you have to take care of the key rotation and decide where to store the SA's secrets/keys to keep it safe, although theoretically it works. B is the most elegant way to solve this issue as GCP client will automatically get the token from instance internally without requiring you to consider the headaches you encounter in option C.

upvoted 2 times

khadar2001 1 year, 4 months ago

This is verified from Practice Exam

upvoted 1 times

hpf97 1 year, 6 months ago

A & D does not deal about connection C : is mandatory for security reason B : could be true but here is always a "default" scope. And the error message is about connection not access. And it is for OAuth2 which we do not know if the python script uses it.

upvoted 2 times

ryaryarya 1 year, 7 months ago

This is obviously a really old question. The best answer is B. Here's why: - There was a time when you couldn't update access scopes on a VM and had to create a new one if the scopes needed to change (which was painful). - Without access scopes enabled, Google will block access to those APIs, so a new service account isn't enough, even if IAM was a problem. - The question just says that the script 'cannot connect' to BQ - it doesn't mention permissions.

upvoted 2 times

gcloud007 1 year, 7 months ago

Actually the correct answer is B, reason is , you need to define the scope first, service accounts comes later.

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

A and C are correct, but we eliminated A because they mentioned "cannot connect" which means the script can run which means the client library was already installed, so final answer is only "C" "C" was chosen because in order to access BigQuery, the script needs to authenticate and be authorized. The recommended way to do this for applications running on Compute Engine is to use a service account. Create a service account with the appropriate permissions (e.g., "BigQuery Data Editor") to access your BigQuery data. When running the script, make sure it uses the service account credentials to authenticate. This can be done by setting the GOOGLE_APPLICATION_CREDENTIALS environment variable to the path of the service account key file.

upvoted 2 times

Hungdv 2 years ago

Choose C

upvoted 1 times

kingfighers 2 years, 2 months ago

I suppose all of them are correct, but we should choose the least effort, B is correct..

upvoted 1 times

kingfighers 2 years, 2 months ago

I suppose all of them are correct, but we should choose the least effort, B is correct..run script on a new vm, not create a new vm..

upvoted 1 times

kingfighers 2 years, 2 months ago

run script on a new vm, not create a new vm..

upvoted 1 times

a2le 2 years, 2 months ago

Tricky question. However, as you can read in gcloud compute instances create documentation: --scopes=[SCOPE,...] If not provided, the instance will be assigned the default scopes, described below. However, if neither --scopes nor --no-scopes are specified and the project has no default service account, then the instance will be created with no scopes. Note that the level of access that a service account has is determined by a combination of access scopes and IAM roles so you must configure both access scopes and IAM roles for the service account to work properly. So, probably, B is the right one, as for the "new vm", I guess that this is because you don't want to stop the current one before having the working one ready...

upvoted 3 times

Robert0 2 years, 2 months ago

C - service account is Google Cloud's best practice

upvoted 1 times

researched_answer_boi 2 years, 3 months ago

You don't need to create a new VM to have different access scopes: https://cloud.google.com/compute/docs/access/service-accounts#accesscopesiam This weakens answer B. When a user-managed service account is attached to the instance, the access scope defaults to cloud-platform: https://cloud.google.com/compute/docs/access/service-accounts#scopes_best_practice See Step 6 in: https://cloud.google.com/compute/docs/instances/change-service-account#changeserviceaccountandscopes These facts leave C as the valid answer.

upvoted 2 times

Question #11 *Topic 1*

Your customer is moving an existing corporate application to Google Cloud Platform from an on-premises data center. The business owners require minimal user disruption. There are strict security team requirements for storing passwords.
What authentication strategy should they use?

A. Use G Suite Password Sync to replicate passwords into Google

B. Federate authentication via SAML 2.0 to the existing Identity Provider

C. Provision users in Google using the Google Cloud Directory Sync tool

D. Ask users to set their Google password to match their corporate password

Suggested Answer: *C*
Provision users to Google's directory The global Directory is available to both Cloud Platform and G Suite resources and can be provisioned by a number of means. Provisioned users can take advantage of rich authentication features including single sign-on (SSO), OAuth, and two-factor verification. You can provision users automatically using one of the following tools and services: Google Cloud Directory Sync (GCDS)

Google Admin SDK -

A third-party connector - GCDS is a connector that can provision users and groups on your behalf for both Cloud Platform and G Suite. Using GCDS, you can automate the addition, modification, and deletion of users, groups, and non-employee contacts. You can synchronize the data from your LDAP directory server to your Cloud Platform domain by using LDAP queries. This synchronization is one-way: the data in your LDAP directory server is never modified. Reference: https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#authentication-and-identity

Community vote distribution

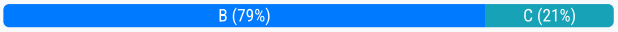

gcp_aws Highly Voted 6 years, 3 months ago
The correct answer is B. GCDS tool only copies the usernames, not the passwords. And more over strict security requirements for the passwords. Not allowed to copy them onto Google, I think. Federation technique help resolve this issue. Please correct me if I am wrong.Passwords are also synchronized: https://support.google.com/a/answer/6120130?hl=en&ref_topic=2679497GCDS synchronises password as well and that is the reason why B is the correct answer. Only in B the password doesn't get copied to GCP.C is the answeryou mistaken GCDS for GSPS, from google site "GSPS won't sync an Active Directory password with a Google Account until it's changed." this if from google to for GCDC "Using GCDS–The recommended way to add users to your Google Account in an Active Directory environment is with Google Cloud Directory Sync (GCDS). GCDS automatically syncs user accounts in your Google domain with user accounts in your Active Directory system."
upvoted 82 times

gcp_aws Highly Voted 6 years, 3 months ago
The correct answer is B. GCDS tool only copies the usernames, not the passwords. And more over strict security requirements for the passwords. Not allowed to copy them onto Google, I think.

Federation technique help resolve this issue. Please correct me if I am wrong.
upvoted 82 times

Neferith 3 years, 11 months ago
Passwords are also synchronized: https://support.google.com/a/answer/6120130?hl=en&ref_topic=2679497
upvoted 8 times

ExamTopicsFan 5 years, 2 months ago
GCDS synchronises password as well and that is the reason why B is the correct answer. Only in B the password doesn't get copied to GCP.
upvoted 11 times

zr79 3 years, 10 months ago
C is the answer
upvoted 3 times

hafid 6 years, 2 months ago

you mistaken GCDS for GSPS, from google site "GSPS won't sync an Active Directory password with a Google Account until it's changed." this if from google to for GCDC "Using GCDS–The recommended way to add users to your Google Account in an Active Directory environment is with Google Cloud Directory Sync (GCDS). GCDS automatically syncs user accounts in your Google domain with user accounts in your Active Directory system."

upvoted 9 times

Eroc 6 years, 10 months ago

"A" will syncronise passwords between on pre-mise and the GCP, this duplicates the existing strategy plus Google's "built-in" encryption of all the data. "B" does not support the moving to GCP. "C" The directory sync tool copies the filesystem settings between servers, UNIX filesystems have permission settings built in and passwords to log into the permission groups, syncing these would set GCP up the same way their on-premises is, plus Google's "built-in" encryption. "D" disrupts the users, so this is not correct. The debate should be between "A" and "C", "C" includes "A" according to (https://cloud.google.com/solutions/migrating-consumer-accounts-to-cloud-identity-or-g-suite-best-practices-federation) so choose "C"B is supported read https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-configuring-single-sign-onThere is no mention SSO is needed.B is ok.miss typed.. C is okB, you dont want to store password as per security guidelines provided in question.GCDS syncs user accounts and some other LDAP attributes but not the passwords, with hybrid connectivity to GCP, SAML (or federation) is the preferred method. Answer should be "B" https://cloud.google.com/solutions/patterns-for-authenticating-corporate-users-in-a-hybrid-environment https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-synchronizing-user-accounts#deciding_what_to_provisionThis is the best answer so far.GCDS does sync passwords. Please refer - https://support.google.com/a/answer/6120130. Since the question says client wants to move to GCP , C should be the answer.In Active Directory, passwords are stored as write-only. They can't be read through any interface, such as LDAP. Therefore, conventional synchronization methods (for example, Google Cloud Directory Sync) can't access them. The only way to read passwords is to capture them when they're set or changed.The article implies that ADFS is best but suggests you also need the GCDS. This make

upvoted 22 times

Eroc 6 years, 10 months ago

"A" will syncronise passwords between on pre-mise and the GCP, this duplicates the existing strategy plus Google's "built-in" encryption of all the data. "B" does not support the moving to GCP. "C" The directory sync tool copies the filesystem settings between servers, UNIX filesystems have permission settings built in and passwords to log into the permission groups, syncing these would set GCP up the same way their on-premises is, plus Google's "built-in" encryption. "D" disrupts the users, so this is not correct. The debate should be between "A" and "C", "C" includes "A" according to (https://cloud.google.com/solutions/migrating-consumer-accounts-to-cloud-identity-or-g-suite-best-practices-federation) so choose "C"

upvoted 22 times

Gobblegobble 6 years, 1 month ago

B is supported read https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-configuring-single-sign-on

upvoted 4 times

Gobblegobble 6 years, 1 month ago

B is supported read https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-configuring-single-sign-onThere is no mention SSO is needed.

upvoted 4 times

tsys 5 years, 5 months ago

There is no mention SSO is needed.

upvoted 3 times

tartar 6 years ago

B is ok.

upvoted 5 times

tartar 6 years ago

B is ok.miss typed.. C is ok

upvoted 5 times

tartar 6 years ago

miss typed.. C is ok

upvoted 11 times

nitinz 5 years, 5 months ago

B, you dont want to store password as per security guidelines provided in question.

upvoted 3 times

cetanx 6 years, 1 month ago

GCDS syncs user accounts and some other LDAP attributes but not the passwords, with hybrid connectivity to GCP, SAML (or federation) is the preferred method.

Answer should be "B"

https://cloud.google.com/solutions/patterns-for-authenticating-corporate-users-in-a-hybrid-environment https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-synchronizing-user-accounts#deciding_what_to_provision

upvoted 16 times

cetanx 6 years, 1 month ago

GCDS syncs user accounts and some other LDAP attributes but not the passwords, with hybrid connectivity to GCP, SAML (or federation) is the preferred method. Answer should be "B" https://cloud.google.com/solutions/patterns-for-authenticating-corporate-users-in-a-hybrid-environment https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-synchronizing-user-accounts#deciding_what_to_provisionThis is the best answer so far.GCDS does sync passwords. Please refer - https://support.google.com/a/answer/6120130. Since the question says client wants to move to GCP , C should be the answer.In Active Directory, passwords are stored as write-only. They can't be read through any interface, such as LDAP. Therefore, conventional synchronization methods (for example, Google Cloud Directory Sync) can't access them. The only way to read passwords is to capture them when they're set or changed.The article implies that ADFS is best but suggests you also need the GCDS. This makes sense, you need the users in Google to allocate permissions but you don't want to copy the passwords across hence ADFS.

upvoted 16 times

squishy_fishy 4 years, 7 months ago

This is the best answer so far.

upvoted 2 times

SamirJ 5 years, 10 months ago

GCDS does sync passwords. Please refer - https://support.google.com/a/answer/6120130. Since the question says client wants to move to GCP , C should be the answer.

upvoted 5 times

SamirJ 5 years, 10 months ago

GCDS does sync passwords. Please refer - https://support.google.com/a/answer/6120130. Since the question says client wants to move to GCP , C should be the answer.In Active Directory, passwords are stored as write-only. They can't be read through any interface, such as LDAP. Therefore, conventional synchronization methods (for example, Google Cloud Directory Sync) can't access them. The only way to read passwords is to capture them when they're set or changed.

upvoted 5 times

MilkyMist 1 year, 7 months ago

In Active Directory, passwords are stored as write-only. They can't be read through any interface, such as LDAP. Therefore, conventional synchronization methods (for example, Google Cloud Directory Sync) can't access them. The only way to read passwords is to capture them when they're set or changed.

upvoted 1 times

BiddlyBdoyng 3 years, 1 month ago

The article implies that ADFS is best but suggests you also need the GCDS. This makes sense, you need the users in Google to allocate permissions but you don't want to copy the passwords across hence ADFS.

upvoted 2 times

Ral17 Most Recent 4 months, 3 weeks ago

The correct answer is B. SAML 2.0 Federation means --> Google redirects to your corporate IDP --> IDP validates credentials --> Sends SAML assertion back to Google --> Access granted.

upvoted 1 times

j_vish 1 year, 5 months ago

B is most correct answer

upvoted 1 times

rahulrawat0789 1 year, 6 months ago

Its B, because of strict security requirements. User will not be forced move passwords to Google cloud

upvoted 2 times

hpf97 1 year, 6 months ago

A&C: would also copy password in GCP, so will certainly not satisfy security requirement B : Recommended, no service interruption, no change for users, quick to set up https://cloud.google.com/architecture/identity/best-practices-for-federating D : Just kidding solution

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Minimal User Disruption: Federated authentication allows users to use their existing corporate credentials to access the application in Google Cloud. This eliminates the need for them to create and remember new passwords, minimizing disruption and improving user experience. 2. Strict Security Requirements: SAML 2.0 is a widely used, secure standard for authentication and authorization. It allows the existing identity provider (IdP) to handle password management and security policies, ensuring compliance with the security team's requirements. 3. Centralized Identity

Management: Federation keeps identity management centralized within the existing corporate infrastructure. This simplifies user management and reduces the overhead of managing identities in multiple places.

upvoted 2 times

selected 1 year, 10 months ago

cross-domain SSO can be achieved by SAML

upvoted 1 times

JohnJamesB1212 1 year, 11 months ago

I think B is correct

upvoted 1 times

maxdanny 1 year, 11 months ago

Minimal user disruption: By federating authentication via SAML 2.0, users can continue using their existing corporate credentials without having to manage or remember new passwords. Security requirements: SAML 2.0 federation allows your organization to maintain control over user authentication and password management within the existing Identity Provider (IdP). Passwords do not need to be stored in Google's systems, which aligns with strict security requirements.

upvoted 1 times

Manishjb006 2 years ago

B is right one . Because C While Google Cloud Directory Syc (GCDS) helps sync users between an on-premises directory and Google, it does not address the password management aspect. Users may still face disruptions as this method might not handle existing passwords securely.

upvoted 1 times

Hungdv 2 years ago

Choose B

upvoted 1 times

kingfighers 2 years, 2 months ago

the most convenient way is B, but the principle of this kind of exam is use cloud provider's native tools, so the C is correct.. this principle is also used on aws

upvoted 1 times

santoshchauhan 2 years, 5 months ago

B. Federate authentication via SAML 2.0 to the existing Identity Provider. Here's why: Security: SAML 2.0 allows for secure single sign-on (SSO) without storing passwords on Google's side. It ensures that authentication happens against the corporate Identity Provider (IdP), which maintains control over the user credentials. Minimal Disruption: Users can continue to use their existing corporate credentials to access the application on GCP without having to remember a new set of credentials or go through a password change process. Compliance: It satisfies the security team's requirements for password storage by ensuring that passwords remain within the corporate boundary. Integration: SAML is widely supported and can be integrated with many IdPs, allowing for a seamless transition to cloud-based resources while leveraging existing identity management infrastructure.

upvoted 6 times

lisabisa 2 years, 6 months ago

The correct answer is C. Google Cloud Directory Sync will provide federated authentications. B is wrong because SAML is used for Single sign-on. It also doesn't mention how the cloud can be authenticated to the existing Identity Provider. SAML by itself is not enough to do the job.

upvoted 2 times

xxoox 2 years, 6 months ago

Federating authentication aligns with strict security team requirements for password storage, as it avoids the need to store or sync passwords outside the corporate environment.

upvoted 2 times

hzaoui 2 years, 7 months ago

Minimal User Disruption: Users continue using their existing corporate credentials for both on-premises and GCP applications, avoiding password resets or new account creations. Security Team Requirements: GCP doesn't store or manage corporate passwords; authentication relies on the existing Identity Provider (IdP), meeting strict password storage requirements.

upvoted 1 times

Question #12 *Topic 1*

Your company has successfully migrated to the cloud and wants to analyze their data stream to optimize operations. They do not have any existing code for this analysis, so they are exploring all their options. These options include a mix of batch and stream processing, as they are running some hourly jobs and live- processing some data as it comes in.
Which technology should they use for this?

A. Google Cloud Dataproc

B. Google Cloud Dataflow

C. Google Container Engine with Bigtable

D. Google Compute Engine with Google BigQuery

Suggested Answer: *B*
Cloud Dataflow is a fully-managed service for transforming and enriching data in stream (real time) and batch (historical) modes with equal reliability and expressiveness -- no more complex workarounds or compromises needed. Reference: https://cloud.google.com/dataflow/

Community vote distribution
B (100%)

Eroc Highly Voted 1 year, 11 months ago
All four options can accomplish what the question asks, in regards to batching and streaming processes. "A" is for Apache Spark and Hadoop, a juggernaut in speed of data processing. "B" is Google's best attempt at TIBCO, Ab Initio, and other processing technology, built explicity for visualizing batch operations and streams without through various labeled circuit boards. "C" and "D" are used within "A" and "B" and would require more work and higher risk. I'd guess Google wants you to select "B"
upvoted 40 times

2g Highly Voted 6 years, 6 months ago
answer: B
upvoted 6 times

Ral17 Most Recent 4 months, 3 weeks ago
Dataflow is built on Apache Beam, which was specifically designed for unified batch + stream processing. Google managed underlying infra and autoscales based on volume.
upvoted 1 times

RamDhangar 7 months, 3 weeks ago
answer:B
upvoted 1 times

hpf97 1 year, 6 months ago
DataFlow is a managed service which processes stream and batch data. So it prefectly fits the need
upvoted 2 times

Ekramy_Elnaggar 1 year, 9 months ago
1. Unified Batch and Stream Processing: Dataflow is a fully managed service designed for both batch and stream data processing. This makes it ideal for your company's needs, as they require both hourly batch jobs and live stream processing. 2. No Existing Code: Dataflow provides a unified programming model and SDKs (Java, Python) for building data pipelines, which is beneficial since your company doesn't have existing code and needs to develop new solutions. 3. Serverless and Scalable: Dataflow is serverless, meaning you don't need to manage infrastructure. It automatically scales resources based on the workload, ensuring efficient processing of both batch and stream data. 4. Cost-Effective: Dataflow's autoscaling and pay-per-use model optimize costs by only utilizing resources when needed.
upvoted 6 times

Singapore123 1 year, 10 months ago
B. Google Cloud Dataflow Explanation: Unified Processing: Google Cloud Dataflow is designed to handle both batch and stream processing in a unified manner. This means you can process data as it arrives (stream processing) and also perform scheduled batch jobs efficiently. Serverless and Scalable: Dataflow is serverless, which means you don't have to worry about managing the underlying infrastructure. It automatically scales to handle varying workloads, making it ideal for optimizing operations based on live data streams and scheduled jobs. Integration with Other Google Cloud Services: Dataflow integrates well with other Google Cloud services, such as Google Cloud Storage, BigQuery, and Pub/Sub. This makes it easier to

build a comprehensive data pipeline that can analyze data streams effectively. Flexible SDKs: Dataflow supports popular programming languages like Java and Python, allowing your team to write custom processing logic as needed.

upvoted 4 times

Hungdv 2 years ago

Choose B

upvoted 1 times

hzaoui 2 years, 7 months ago

B is correct

upvoted 1 times

devakram 2 years, 8 months ago

chatGPT answers: B. Google Cloud Dataflow Google Cloud Dataflow is a fully managed service for stream and batch data processing. It is built on Apache Beam and provides a unified programming model, making it an ideal choice for scenarios where both batch and stream data processing are required. Dataflow simplifies the complexities of data parallel processing, allowing for easy development and maintenance of data processing pipelines. It integrates well with other Google Cloud services, like BigQuery for analytics and Cloud Storage for storing data, providing a comprehensive solution for real-time and batch data processing needs.

upvoted 2 times

BiddlyBdoyng 3 years, 1 month ago

The word analysis throws me off. Wonder if the question is just written incorrectly here? I'd say Dataflow is a key tool to enable the processing of the data to be able to do the analysis but feels like the final analysis should be in a database.

upvoted 5 times

alekonko 3 years, 4 months ago

B is the answer

upvoted 2 times

Deb2293 3 years, 5 months ago

A is a managed Hadoop and Spark service. C and D are mostly for petabyte kinds of data. So remains B (suitable for ETL jobs)

upvoted 2 times

omermahgoub 3 years, 8 months ago

To analyze a data stream and optimize operations, your company could consider using Google Cloud Dataflow, which is a fully-managed, cloud-native data processing service that can handle both batch and stream processing. Google Cloud Dataflow is designed to handle large volumes of data and can scale up or down automatically to meet the needs of the workload. It provides a number of pre-built connectors and integrations that make it easy to ingest data from a variety of sources, and it offers a range of processing options, including batch processing and stream processing, that can be used to analyze the data in real-time. Option A: Google Cloud Dataproc, option C: Google Container Engine with Bigtable, and option D: Google Compute Engine with Google BigQuery, while potentially useful for certain types of data processing, would not necessarily be well-suited to handle both batch and stream processing in the way that Google Cloud Dataflow can

upvoted 3 times

thamaster 3 years, 8 months ago

answer is D for me the question is which tool for analyse data. Dataflow does not analyse data

upvoted 2 times

megumin 3 years, 9 months ago

ok for B

upvoted 1 times

Mahmoud_E 3 years, 10 months ago

B is the right answer

upvoted 1 times

Question #13 *Topic 1*

Your customer is receiving reports that their recently updated Google App Engine application is taking approximately 30 seconds to load for some of their users.
This behavior was not reported before the update.
What strategy should you take?

A. Work with your ISP to diagnose the problem

B. Open a support ticket to ask for network capture and flow data to diagnose the problem, then roll back your application

C. Roll back to an earlier known good release initially, then use Stackdriver Trace and Logging to diagnose the problem in a development/test/staging environment

D. Roll back to an earlier known good release, then push the release again at a quieter period to investigate. Then use Stackdriver Trace and Logging to diagnose the problem

Suggested Answer: *C*
Stackdriver Logging allows you to store, search, analyze, monitor, and alert on log data and events from Google Cloud Platform and Amazon Web Services (AWS). Our API also allows ingestion of any custom log data from any source. Stackdriver Logging is a fully managed service that performs at scale and can ingest application and system log data from thousands of VMs. Even better, you can analyze all that log data in real time. Reference: https://cloud.google.com/logging/

Community vote distribution
C (93%)

TosO Highly Voted 6 years, 9 months ago
C is the answer
upvoted 28 times

MyPractice Highly Voted 1 year, 11 months ago
Key word: This behavior was not reported before the update A - Not Correct as it was working before with same ISP B - New code update caused an issue- why to open support ticket C - I agree with C D - This requires downtime and live prod affected too
upvoted 20 times

MyPractice 1 year, 11 months ago
Key word: This behavior was not reported before the update A - Not Correct as it was working before with same ISP B - New code update caused an issue- why to open support ticket C - I agree with C D - This requires downtime and live prod affected too"then use Stackdriver Trace and Logging to diagnose the problem in a development/test/staging environment" they are NOT asking us to setup Dev/Text/Stage.. meaning the environment already exist and we have to use it"then use Stackdriver Trace and Logging to diagnose the problem in a development/test/staging environment" this is not asking for set environment either, it just says to diagnose problem in other environment so C it is
upvoted 20 times

MyPractice 6 years, 7 months ago
"then use Stackdriver Trace and Logging to diagnose the problem in a development/test/staging environment" they are NOT asking us to setup Dev/Text/Stage.. meaning the environment already exist and we have to use it"then use Stackdriver Trace and Logging to diagnose the problem in a development/test/staging environment" this is not asking for set environment either, it just says to diagnose problem in other environment so C it is
upvoted 1 times

MyPractice 6 years, 7 months ago
"then use Stackdriver Trace and Logging to diagnose the problem in a development/test/staging environment" they are NOT asking us to setup Dev/Text/Stage.. meaning the environment already exist and we have to use it
upvoted 1 times

hafid 6 years, 2 months ago
"then use Stackdriver Trace and Logging to diagnose the problem in a development/test/staging environment" this is not asking for set environment either, it just says to diagnose problem in other environment so C it is
upvoted 1 times

WayneHsu Most Recent 3 months, 3 weeks ago
The answer is C.
upvoted 1 times

gnsridhar 7 months ago

D is more appropriate as the assumption is that the version deployed is already tested on dev/test/stg, which rules out option C

upvoted 1 times

Upender_PDE 8 months, 1 week ago

Answer: C Best practice incident response: Restore service immediately - Roll back to the known good version Investigate safely - Use Stackdriver Trace and Logging in non-production environments Fix and test thoroughly - Before redeploying to production This minimizes user impact while still allowing you to properly diagnose and fix the root cause.

upvoted 1 times

Qilee 1 year, 3 months ago

C provides all options to provide and backup this situation.

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Prioritize User Experience: Rolling back to a stable version quickly minimizes user impact and restores the application to a functional state. This should be the immediate first step. 2. Controlled Environment: Diagnosing the issue in a development/test/staging environment allows you to investigate without affecting real users. You can reproduce the problem, gather data, and test potential solutions safely. 3. Powerful Diagnostic Tools: Stackdriver Trace helps you pinpoint performance bottlenecks by tracing requests across your application. Stackdriver Logging provides detailed logs to understand application behavior and identify errors.

upvoted 5 times

hzaoui 2 years, 7 months ago

C is correct

upvoted 1 times

AdityaGupta 2 years, 10 months ago

Your customer is receiving reports that their recently updated Google App Engine application is taking approximately 30 seconds to load for some of their users. This behavior was not reported before the update. What strategy should you take? Here the application (our code) is updated and only some users are facing lantecy (Cloud Trace) issue. The issue is not with ISP (A), Not an issue with Google (B). Rollback must be done as mitigation, but testing should be done in Non-Prod environments (C), not on prod environment (D). Hence C is correct answer.

upvoted 1 times

jrisl1991 2 years, 10 months ago

I'm going for C. While D may be "better" in case this is an issue that only occurs in production, I think that keeping the disruption at minimum would be the best practice, which D would not really do. Plus, if the problem is load related, having this released at a quieter period may not surface the problem either.

upvoted 2 times

frankryuu 3 years, 1 month ago

Although it sounds like the right answer to do network tracing in stg again, this may be a network pass-through related issue and it is felt that the problem may not be reproduced if not checked in a prod environment.

upvoted 2 times

FigVam 3 years, 3 months ago

should be C

upvoted 2 times

alekonko 3 years, 4 months ago

C is the answer

upvoted 2 times

JC0926 3 years, 5 months ago

Option C is also a valid strategy in this scenario. Rolling back to an earlier known good release initially and using Stackdriver Trace and Logging to diagnose the problem in a development/test/staging environment can help diagnose the issue without impacting production users. However, the reason why option D may be a better approach is that it allows for investigation during a quieter period, which can reduce the impact of any issues that may occur during the investigation. Rolling back to a known good release and then pushing the release again at a quieter period can help to ensure that users are not impacted during the investigation.

upvoted 3 times

megumin 3 years, 9 months ago

ok for C

upvoted 1 times

Mahmoud_E 3 years, 10 months ago

C is the correct answer

upvoted 1 times

Question #14 *Topic 1*

A production database virtual machine on Google Compute Engine has an ext4-formatted persistent disk for data files. The database is about to run out of storage space.
How can you remediate the problem with the least amount of downtime?

A. In the Cloud Platform Console, increase the size of the persistent disk and use the resize2fs command in Linux.

B. Shut down the virtual machine, use the Cloud Platform Console to increase the persistent disk size, then restart the virtual machine

C. In the Cloud Platform Console, increase the size of the persistent disk and verify the new space is ready to use with the fdisk command in Linux

D. In the Cloud Platform Console, create a new persistent disk attached to the virtual machine, format and mount it, and configure the database service to move the files to the new disk

E. In the Cloud Platform Console, create a snapshot of the persistent disk restore the snapshot to a new larger disk, unmount the old disk, mount the new disk and restart the database service

Suggested Answer: *A*
On Linux instances, connect to your instance and manually resize your partitions and file systems to use the additional disk space that you added. Extend the file system on the disk or the partition to use the added space. If you grew a partition on your disk, specify the partition. If your disk does not have a partition table, specify only the disk ID. sudo resize2fs /dev/[DISK_ID][PARTITION_NUMBER] where [DISK_ID] is the device name and [PARTITION_NUMBER] is the partition number for the device where you are resizing the file system. Reference: https://cloud.google.com/compute/docs/disks/add-persistent-disk

Community vote distribution
A (91%) (3

TosO Highly Voted 6 years, 9 months ago
A is the correct answer because the question says "with minimum downtime"
upvoted 28 times

passnow Highly Voted 6 years, 8 months ago
least amount of downtime? is the sugar word. You miss that you miss all. Everything there is correct but I believe its only A that fits that requirementbut in option A, nowhere it is mentioned to shut down the VM.No need to reboot.
upvoted 14 times

passnow 6 years, 8 months ago
least amount of downtime? is the sugar word. You miss that you miss all. Everything there is correct but I believe its only A that fits that requirement
upvoted 14 times

raj117 5 years, 4 months ago
but in option A, nowhere it is mentioned to shut down the VM.No need to reboot.
upvoted 1 times

raj117 5 years, 4 months ago
but in option A, nowhere it is mentioned to shut down the VM.
upvoted 1 times

monkeym 5 years, 1 month ago
No need to reboot.
upvoted 1 times

Upender_PDE Most Recent 8 months, 1 week ago
Answer: A Process: In Cloud Console: Increase the persistent disk size (done online, no downtime) SSH into the VM: Run sudo resize2fs /dev/[DEVICE_NAME] to expand the filesystem Verify: Check with df -h that the new space is available
upvoted 1 times

JoseCloudEng1994 11 months, 3 weeks ago
The answer is A. In real life I have done this, you can just edit the instance disk on GCP. Though I have asked for the instance to be restarted to make sure its all good so learning this command is good too
upvoted 2 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Online Resizing: Google Cloud Platform allows you to increase the size of a persistent disk while it's attached to a running VM. This means you don't need to shut down the database server. 2. resize2fs: This Linux command extends the file system to utilize the newly added space on the disk. It can be run while the file system is mounted, minimizing downtime.

upvoted 4 times

SerGCP 1 year, 9 months ago

Selected Answer: C

A) can not work becouse you must extend partition and after than you can extend filesystem C) ok becouse you can use fdisk in linux to extend partition and after than you can resize ext4 filesystem

upvoted 1 times

SerGCP 1 year, 9 months ago

A) can not work becouse you must extend partition and after than you can extend filesystem C) ok becouse you can use fdisk in linux to extend partition and after than you can resize ext4 filesystemWhile fdisk can show the partition table, it cannot dynamically expand the filesystem to use the new space. You still need a command like resize2fs (for ext4) or xfs_growfs (for XFS).

upvoted 1 times

ast3citos 10 months, 1 week ago

While fdisk can show the partition table, it cannot dynamically expand the filesystem to use the new space. You still need a command like resize2fs (for ext4) or xfs_growfs (for XFS).

upvoted 1 times

Nora9 2 years, 9 months ago

A is the right answer.You should resize the disk, take a snapshot, then resize the filesystem and partitions (eg.) ext4, xfs etc.

upvoted 1 times

AdityaGupta 2 years, 10 months ago

Unlike Azure, in google you can dynamically resize the persistent disk while VM is running. This narrows down the option to A or C. Since the question says "ext4-formatted persistent disk", we need to choose correct command (resize2fs or fdisk) for Linux for resizing ext4 file format disk. To resize an ext4 file system in Linux, you can use the resize2fs command. FDISK to manipulate partition tables in Linux.

upvoted 8 times

blackhawk86 2 years, 11 months ago

According to the right url, A is the right answer. https://cloud.google.com/compute/docs/disks/resize-persistent-disk

upvoted 4 times

blackhawk86 2 years, 11 months ago

The right URL for the oficial document is, https://cloud.google.com/compute/docs/disks/resize-persistent-disk

upvoted 1 times

heretolearnazure 2 years, 12 months ago

A is no brainer

upvoted 1 times

jalberto 3 years ago

E is the correct because is true you need minimum downtime but in Production a backup is a must.

upvoted 2 times

jalberto 3 years ago

E because you are in Production, and you need a backup

upvoted 1 times

alekonko 3 years, 4 months ago

A is correct, resize disk don't required reboot or downtime https://cloud.google.com/compute/docs/disks/resize-persistent-disk

upvoted 3 times

omermahgoub 3 years, 8 months ago

A: Increasing the size of the persistent disk in the Cloud Platform Console and using the resize2fs command in Linux. Increasing the size of the persistent disk can be done without requiring the virtual machine to be shut down, and the resize2fs command can be used to resize the ext4 filesystem on the disk to take advantage of the additional space. This will allow you to add more storage space to the virtual machine without disrupting the database service.

upvoted 1 times

megumin 3 years, 9 months ago

A is ok

upvoted 2 times

Mahmoud_E 3 years, 10 months ago

A https://cloud.google.com/compute/docs/disks/resize-persistent-disk?_ga=2.233866652.-3622898.1631303718

upvoted 1 times

Question #15 *Topic 1*

Your application needs to process credit card transactions. You want the smallest scope of Payment Card Industry (PCI) compliance without compromising the ability to analyze transactional data and trends relating to which payment methods are used.
How should you design your architecture?

A. Create a tokenizer service and store only tokenized data

B. Create separate projects that only process credit card data

C. Create separate subnetworks and isolate the components that process credit card data

D. Streamline the audit discovery phase by labeling all of the virtual machines (VMs) that process PCI data

E. Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor

Suggested Answer: *A*
Reference: https://www.sans.org/reading-room/whitepapers/compliance/ways-reduce-pci-dss-audit-scope-tokenizing-cardholder-data-33194

Community vote distribution
A (100%)

AD2AD4 Highly Voted 6 years, 2 months ago
Final Decision to go with Option A. I have done PCI DSS Audit for my project and thats the best suited case. 100% sure to use tokenised data instead of actual card numberBut with A you cannot extract statistics. That is the second r4equirement.Thinking about that better, I think you can because you are only tokenizing the sensitive data, not the transaction type.You can as the generated token for a given credit card would be same(generally but there are approaches which can give you different token for the same sensitive data input). Only thing that you won't know is the actual card number which is not required for the trend analysis. When the trend analysis involves referential integrity then tokenization process becomes challenging but still once data is tokenized correctly you should be able to perform any kind of the analysis.Analyzing Transaction does not require Credit Card number I guess. Only amount of transaction or balance what is needed. We also perform something similar with transactional data with tokenized PII information. So CC can be tokenized. So answer should be A.I agree. A is the best option
upvoted 48 times

AD2AD4 Highly Voted 6 years, 2 months ago
Final Decision to go with Option A. I have done PCI DSS Audit for my project and thats the best suited case. 100% sure to use tokenised data instead of actual card number
upvoted 48 times

Musk 6 years, 2 months ago
But with A you cannot extract statistics. That is the second r4equirement.
upvoted 4 times

Musk 6 years, 2 months ago
But with A you cannot extract statistics. That is the second r4equirement.Thinking about that better, I think you can because you are only tokenizing the sensitive data, not the transaction type.You can as the generated token for a given credit card would be same(generally but there are approaches which can give you different token for the same sensitive data input). Only thing that you won't know is the actual card number which is not required for the trend analysis. When the trend analysis involves referential integrity then tokenization process becomes challenging but still once data is tokenized correctly you should be able to perform any kind of the analysis.Analyzing Transaction does not require Credit Card number I guess. Only amount of transaction or balance what is needed. We also perform something similar with transactional data with tokenized PII information. So CC can be tokenized. So answer should be A.
upvoted 4 times

Musk 6 years ago
Thinking about that better, I think you can because you are only tokenizing the sensitive data, not the transaction type.
upvoted 2 times

RitwickKumar 4 years ago
You can as the generated token for a given credit card would be same(generally but there are approaches which can give you different token for the same sensitive data input). Only thing that you won't know is the actual card number which is not required for the trend analysis. When the trend analysis involves referential integrity then tokenization process becomes challenging but still once data is tokenized correctly you should be able to perform any kind of the analysis.
upvoted 2 times

Arimaverick 5 years, 7 months ago

Analyzing Transaction does not require Credit Card number I guess. Only amount of transaction or balance what is needed. We also perform something similar with transactional data with tokenized PII information. So CC can be tokenized. So answer should be A.

upvoted 6 times

AzureDP900 3 years, 10 months ago

I agree. A is the best option

upvoted 3 times

omermahgoub 3 years, 8 months ago

To minimize the scope of Payment Card Industry (PCI) compliance while still allowing for the analysis of transactional data and trends related to payment methods, you should consider using a tokenizer service and storing only tokenized data, as described in option A.

Tokenization is a process of replacing sensitive data, such as credit card numbers, with unique, randomly-generated tokens that cannot be used for fraudulent purposes. By using a tokenizer service and storing only tokenized data, you can reduce the scope of PCI compliance to only the tokenization service, rather than the entire application. This can help minimize the amount of sensitive data that needs to be protected and reduce the overall compliance burden.

upvoted 41 times

omermahgoub 3 years, 8 months ago

To minimize the scope of Payment Card Industry (PCI) compliance while still allowing for the analysis of transactional data and trends related to payment methods, you should consider using a tokenizer service and storing only tokenized data, as described in option A. Tokenization is a process of replacing sensitive data, such as credit card numbers, with unique, randomly-generated tokens that cannot be used for fraudulent purposes. By using a tokenizer service and storing only tokenized data, you can reduce the scope of PCI compliance to only the tokenization service, rather than the entire application. This can help minimize the amount of sensitive data that needs to be protected and reduce the overall compliance burden.man, this is an amazing answer. propsthanks to chagpt, are you serious saying that?Nicely explained, thanks.

upvoted 41 times

oxfordcommaa 3 years, 6 months ago

man, this is an amazing answer. propsthanks to chagpt, are you serious saying that?

upvoted 5 times

oxfordcommaa 3 years, 6 months ago

man, this is an amazing answer. props

upvoted 5 times

ccpmad 2 years, 2 months ago

thanks to chagpt, are you serious saying that?

upvoted 1 times

Saxena_Vibhor 2 years, 7 months ago

Nicely explained, thanks.

upvoted 1 times

Ekramy_Elnaggar Most Recent 1 year, 9 months ago

1. Reduced PCI Scope: Tokenization replaces sensitive credit card data with non-sensitive tokens. This significantly reduces the scope of PCI DSS compliance, as you no longer store actual cardholder data. Only the tokenization service needs to be PCI compliant. 2. Data Analysis: You can still analyze transactional data and trends using the tokenized data. The tokens can be linked back to the original cardholder data if needed for specific analysis or reporting purposes, but this would be done in a controlled and secure environment.

upvoted 2 times

hzaoui 2 years, 7 months ago

A is correct

upvoted 1 times

devinss 2 years, 11 months ago

Not sure why 100% agree on A. To limit PCI DSS scope, the data handling should be done in a separate project with very limited access. Only in this project should tokenization be done and made available for analytics. The first requirement however, is isolating the payment and tokenization code in a separate project. Answer should be B.

upvoted 2 times

heretolearnazure 2 years, 12 months ago

Tokenizing is the best way to protect PCI information.

upvoted 1 times

nocrush 3 years ago

A is my best option

upvoted 1 times

KjChen 3 years, 9 months ago

https://cloud.google.com/architecture/tokenizing-sensitive-cardholder-data-for-pci-dss

upvoted 4 times

andreavale 3 years, 9 months ago

ok for A

upvoted 1 times

minmin2020 3 years, 10 months ago

A. Create a tokenizer service and store only tokenized data

upvoted 1 times

BiddlyBdoyng 3 years, 10 months ago

B appears the most thorough but the question asks to comply with the smallest scope, network segmentation is not a must. Tokenization is simpler. C is similar to B, more than required. D & E do not address the problem.

upvoted 1 times

holerina 3 years, 11 months ago

correct answer is A use tokenize

upvoted 2 times

abirroy 3 years, 11 months ago

Correct answer A

upvoted 1 times

Nirca 4 years, 4 months ago

The mandatory stage in PCI is having a encryption/ description system. Data must not be stored as is with PAN. So A IS A MUST. The rest are nice to have.

upvoted 1 times

ryzior 4 years, 4 months ago

I think it should be A and C - the paper states clearly, a proper network segmentation is still required to disparate the vault and token servers from the rest of the flat network.

upvoted 1 times

sjmsummer 4 years, 7 months ago

I chose A. But why C is not good?

upvoted 1 times

vincy2202 4 years, 7 months ago

A is the correct answer https://cloud.google.com/architecture/tokenizing-sensitive-cardholder-data-for-pci-dss#a_service_for_handling_sensitive_information

upvoted 3 times

Question #16 *Topic 1*

You have been asked to select the storage system for the click-data of your company's large portfolio of websites. This data is streamed in from a custom website analytics package at a typical rate of 6,000 clicks per minute. With bursts of up to 8,500 clicks per second. It must have been stored for future analysis by your data science and user experience teams.
Which storage infrastructure should you choose?

A. Google Cloud SQL

B. Google Cloud Bigtable

C. Google Cloud Storage

D. Google Cloud Datastore

Suggested Answer: *B*
Google Cloud Bigtable is a scalable, fully-managed NoSQL wide-column database that is suitable for both real-time access and analytics workloads. Good for: ➪ Low-latency read/write access ➪ High-throughput analytics ➪ Native time series support Common workloads: ➪ IoT, finance, adtech ➪ Personalization, recommendations ➪ Monitoring ➪ Geospatial datasets ➪ Graphs Incorrect Answers: C: Google Cloud Storage is a scalable, fully-managed, highly reliable, and cost-efficient object / blob store. Is good for: ➪ Images, pictures, and videos ➪ Objects and blobs ➪ Unstructured data D: Google Cloud Datastore is a scalable, fully-managed NoSQL document database for your web and mobile applications. Is good for: ➪ Semi-structured application data ➪ Hierarchical data ➪ Durable key-value data ➪ Common workloads: ➪ User profiles ➪ Product catalogs ➪ Game state Reference: https://cloud.google.com/storage-options/

Community vote distribution



victory108 Highly Voted 5 years, 1 month ago
B. Google Cloud Bigtable
upvoted 12 times

jeff001 Highly Voted 5 years, 4 months ago
B, Bigtable due to the IoT like requirements
upvoted 10 times

shallah Most Recent 1 month ago
Bigtable supports sub millisecond multiple concurrent writes and is able to scale for the load
upvoted 1 times

fff2e69 1 year, 9 months ago
Google Cloud Bigtable is well-suited for handling high-throughput, low-latency workloads like clickstream data. It is optimized for analytics on time-series and event data at scale, and it supports high write rates, with the capacity to handle thousands of writes per second. This makes it ideal for storing large volumes of clickstream data with bursts, ensuring data is available for analysis by data science and user experience teams.
upvoted 4 times

Ekramy_Elnaggar 1 year, 9 months ago
1. High throughput for clickstream data: Bigtable is a NoSQL database designed for high write throughput, making it ideal for handling the continuous stream of click data with bursts up to 8,500 clicks per second. 2. Scalability: Bigtable is highly scalable, allowing you to handle increasing data volumes as your website portfolio grows. 3. Low latency: Bigtable provides low latency data access, which is important for real-time analysis and reporting on clickstream data. 4. Integration with BigQuery: Bigtable integrates well with BigQuery, enabling your data science team to perform complex analysis and generate insights from the clickstream data.
upvoted 4 times

upliftinghut 2 years, 2 months ago
Big table if it needs real time but here the need for analysis is not urgent => Google cloud storage.
upvoted 3 times

pakilodi 2 years, 8 months ago
The right answer is BigTable for a such large volume of data.
upvoted 1 times

Nora9 2 years, 9 months ago
B. Google cloud Bigtable is most suitable as it is for analysis and streamed data.

upvoted 1 times

MartinFish 3 years ago

Note: DataStore is now rename as FireStore.

upvoted 1 times

hiromi 3 years, 4 months ago

B is correct

upvoted 1 times

alekonko 3 years, 4 months ago

Bigtable is a high-perf NoSQL db service that handle large volumes of structured data with low latency

upvoted 3 times

zerg0 3 years, 6 months ago

The Google Cloud Bigtable goes together with the BigQuery. The question itself gives away a bit.

upvoted 1 times

omermahgoub 3 years, 8 months ago

For storing click-data that is streamed in at a rate of 6,000 clicks per minute, with bursts of up to 8,500 clicks per second, and that needs to be stored for future analysis by your data science and user experience teams, you should consider using a scalable, high-performance, and low-latency NoSQL database such as Google Cloud Bigtable, option B. Google Cloud Bigtable is a fully managed, high-performance NoSQL database service that is designed to handle large volumes of structured data with low latency. It is well-suited for storing high-velocity data streams and can scale to handle millions of reads and writes per second. Option A: Google Cloud SQL, option C: Google Cloud Storage, and option D: Google Cloud Datastore, would not be suitable for this use case, as they are not designed to handle high-velocity data streams at this scale.

upvoted 4 times

i_am_robot 3 years, 8 months ago

B. Google Cloud Bigtable Google Cloud Bigtable is a scalable, high-performance NoSQL database that is well-suited for storing large amounts of data with low latency. It is designed for high-throughput workloads such as streaming data, and is able to handle bursts of up to millions of reads and writes per second. Given the high volume of click data that needs to be stored and the requirement for low latency, Google Cloud Bigtable would be a good choice for storing the data. It is able to handle the high rate of incoming data and provide fast access to the data for analysis by the data science and user experience teams. Google Cloud SQL is a fully-managed relational database service, and may not be the best choice for storing high-volume streaming data. Google Cloud Storage is an object storage service, and may not provide the necessary performance for storing and querying large amounts of data in real-time. Google Cloud Datastore is a NoSQL document database, and while it may be suitable for storing large amounts of data, it may not provide the necessary performance for handling high volumes of streaming data.

upvoted 2 times

Melampos 3 years, 8 months ago

https://cloud.google.com/bigtable#section-9

upvoted 1 times

Bry_040706 3 years, 8 months ago

B. Google Cloud Bigtable

upvoted 1 times

Mahmoud_E 3 years, 10 months ago

D is right answer

upvoted 1 times

AniketD 3 years, 9 months ago

Cloud Bigtable has all the features to fulfill the requirements mentioned in the question

upvoted 1 times

Question #17 *Topic 1*

You are creating a solution to remove backup files older than 90 days from your backup Cloud Storage bucket. You want to optimize ongoing Cloud Storage spend.
What should you do?

A. Write a lifecycle management rule in XML and push it to the bucket with gsutil

B. Write a lifecycle management rule in JSON and push it to the bucket with gsutil

C. Schedule a cron script using gsutil ls ɜ€"lr gs://backups/** to find and remove items older than 90 days

D. Schedule a cron script using gsutil ls ɜ€"l gs://backups/** to find and remove items older than 90 days and schedule it with cron

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (100%)

Eroc Highly Voted 6 years, 10 months ago
All four are correct answers. Google has built in cron job schduling with Cloud Schedule, so that would place "D" behind "C" in Google's perspective. Google also has it's own lifecycle management command line prompt gcloud lifecycle so "A" or "B" could be used. JSON is slightly faster than XML because of the "{" verse " " distinguisher, with a Trie tree used for alphanumeric parsing. So between "A" and "B", choose "B". Between "B" and "A", "B" is slightly more efficient from the GCP operator perspective. So choose "B".B is okB is correct. Policy = JSON format. No matter if its AWS or GCP.gsutil command takes only json as input for lifecycle management. In case of API, both XML and json can be used.
https://cloud.google.com/storage/docs/gsutil/commands/lifecycle https://cloud.google.com/storage/docs/xml-api/put-bucket-lifecycle
https://cloud.google.com/storage/docs/json_api/v1/buckets/update
upvoted 37 times

Eroc Highly Voted 6 years, 10 months ago
All four are correct answers. Google has built in cron job schduling with Cloud Schedule, so that would place "D" behind "C" in Google's perspective. Google also has it's own lifecycle management command line prompt gcloud lifecycle so "A" or "B" could be used. JSON is slightly faster than XML because of the "{" verse " " distinguisher, with a Trie tree used for alphanumeric parsing. So between "A" and "B", choose "B". Between "B" and "A", "B" is slightly more efficient from the GCP operator perspective. So choose "B".
upvoted 37 times

tartar 6 years ago
B is ok
upvoted 8 times

nitinz 5 years, 5 months ago
B is correct. Policy = JSON format. No matter if its AWS or GCP.
upvoted 11 times

ghitesh 6 years, 7 months ago
gsutil command takes only json as input for lifecycle management. In case of API, both XML and json can be used.
https://cloud.google.com/storage/docs/gsutil/commands/lifecycle https://cloud.google.com/storage/docs/xml-api/put-bucket-lifecycle
https://cloud.google.com/storage/docs/json_api/v1/buckets/update
upvoted 32 times

clouddude 6 years, 3 months ago
I'll go with B. A is not reasonable because life cycle policies are not written in XML. B is reasonable and is cloud native. C requires a cron script which needs something to run the script and is a non-cloud native approach. D requires a cron script which needs something to run the script and is a non-cloud native approach.
upvoted 17 times

Mitthugcpguru Most Recent 1 year, 8 months ago
Gsutil take only json as input so B answer
upvoted 2 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Lifecycle Management: Google Cloud Storage offers built-in lifecycle management rules specifically designed for automated data retention and deletion. This is the most efficient and cost-effective way to manage your backup files. 2. JSON Format: Lifecycle rules are defined in JSON format. 3. gsutil: The gsutil command-line tool is used to interact with Cloud Storage, including setting lifecycle configuration.

upvoted 1 times

Hungdv 2 years ago

choose B

upvoted 1 times

squishy_fishy 2 years, 9 months ago

The correct available answer is B. But in real life, we use Terraform tfvars file.

upvoted 1 times

eka_nostra 3 years ago

Cloud Storage has lifecycle management rules and could be applied with gsutil and gcloud storage buckets. It is common to use JSON for transferring data.

upvoted 2 times

alekonko 3 years, 4 months ago

B, gsutil can set policy using json file https://cloud.google.com/storage/docs/gsutil/commands/lifecycle#examples

upvoted 1 times

omermahgoub 3 years, 8 months ago

To remove backup files older than 90 days from a Cloud Storage bucket and optimize ongoing Cloud Storage spend, you should consider writing a lifecycle management rule in JSON and pushing it to the bucket with gsutil, as described in option B. Lifecycle management rules allow you to automatically delete objects from a Cloud Storage bucket based on age or other criteria, such as the object's storage class. By writing a rule in JSON and pushing it to the bucket with gsutil, you can specify that objects older than 90 days should be deleted, ensuring that the bucket only contains current backup files and minimizing Cloud Storage spend. Option A, C and D would not be suitable for this use case, as they do not allow you to specify lifecycle management rules that delete objects based on age.

upvoted 3 times

i_am_robot 3 years, 8 months ago

B. Write a lifecycle management rule in JSON and push it to the bucket with gsutil To remove backup files older than 90 days from a Cloud Storage bucket, you can use the lifecycle management feature in Cloud Storage. This feature allows you to specify rules to automatically delete objects based on their age.

upvoted 1 times

Melampos 3 years, 8 months ago

https://cloud.google.com/storage/docs/gsutil/commands/lifecycle

upvoted 1 times

Bry_040706 3 years, 8 months ago

B. Life cycle management using JSON.

upvoted 1 times

AniketD 3 years, 9 months ago

B with JSON option is correct

upvoted 1 times

megumin 3 years, 9 months ago

B is ok

upvoted 1 times

Mahmoud_E 3 years, 10 months ago

B is the right answer

upvoted 1 times

zr79 3 years, 10 months ago

life cycle management is the answer written in JSON format. JSON is easier to write and read compared to XML which you can not use in commands

upvoted 1 times

AzureDP900 3 years, 10 months ago

Lifecycle management with JSON is right .. I will go with B

upvoted 1 times

Question #18 *Topic 1*

Your company is forecasting a sharp increase in the number and size of Apache Spark and Hadoop jobs being run on your local datacenter. You want to utilize the cloud to help you scale this upcoming demand with the least amount of operations work and code change.
Which product should you use?

A. Google Cloud Dataflow

B. Google Cloud Dataproc

C. Google Compute Engine

D. Google Kubernetes Engine

Suggested Answer: *B*
Google Cloud Dataproc is a fast, easy-to-use, low-cost and fully managed service that lets you run the Apache Spark and Apache Hadoop ecosystem on Google Cloud Platform. Cloud Dataproc provisions big or small clusters rapidly, supports many popular job types, and is integrated with other Google Cloud Platform services, such as Google Cloud Storage and Stackdriver Logging, thus helping you reduce TCO.
Reference: https://cloud.google.com/dataproc/docs/resources/faq

Community vote distribution
B (100%)

AWS56 Highly Voted 6 years, 7 months ago
"B. Google Cloud Dataproc" is the answer
upvoted 19 times

VinayakBudapanahalli Highly Voted 5 years, 7 months ago
Dataproc is a managed Spark and Hadoop service that lets you take advantage of open source data tools for batch processing, querying, streaming, and machine learning. Dataproc automation helps you create clusters quickly, manage them easily, and save money by turning clusters off when you don't need them. With less time and money spent on administration, you can focus on your jobs and your data.
https://cloud.google.com/dataproc/docs/concepts/overview#:~:text=Dataproc%20is%20a%20managed%20Spark,%2C%20streaming%2C%20and%20machine%20learning.&text=With%20less%20time%20and%20money,your%20jobs%20and%20your%20data.
upvoted 16 times

VinayakBudapanahalli 5 years, 7 months ago
Dataproc is a managed Spark and Hadoop service that lets you take advantage of open source data tools for batch processing, querying, streaming, and machine learning. Dataproc automation helps you create clusters quickly, manage them easily, and save money by turning clusters off when you don't need them. With less time and money spent on administration, you can focus on your jobs and your data.
https://cloud.google.com/dataproc/docs/concepts/overview#:~:text=Dataproc%20is%20a%20managed%20Spark,%2C%20streaming%2C%20and%20machine%20learning.&text=With%20less%20time%20and%20money,your%20jobs%20and%20your%20data.Agreed
upvoted 16 times

AzureDP900 3 years, 10 months ago
Agreed
upvoted 1 times

Ekramy_Elnaggar Most Recent 1 year, 9 months ago
1. Managed Hadoop and Spark: Dataproc is specifically designed for running and managing Apache Spark and Hadoop clusters, which directly addresses your company's needs. 2. Scalability: Dataproc allows you to easily scale your clusters to handle the increasing number and size of jobs. You can add or remove nodes as needed to accommodate the workload. 3. Minimal Operations Work: Dataproc automates cluster creation, configuration, and management, minimizing the operational overhead. This is crucial since you want to reduce operations work. 4. Code Compatibility: Dataproc is compatible with existing Spark and Hadoop code, so you can migrate your jobs with minimal or no code changes.
upvoted 2 times

JohnJamesB1212 1 year, 11 months ago
B is correct because Dataproc is uded for Apache Hadoop and Spark
upvoted 1 times

eka_nostra 3 years ago
Dataflow for data stream and batch. Dataproc for data process with Apache Spark and Hadoop. Compute Engine for VM. Kubernetes Engine for Kubernetes Cluster with Compute Engine under the hood.

upvoted 3 times

alekonko 3 years, 4 months ago

B, Dataproc is Hadoop/Spark managed service in GCP

upvoted 2 times

examch 3 years, 7 months ago

Dataproc is a fully managed and highly scalable service for running Apache Hadoop, Apache Spark, Apache Flink, Presto, and 30+ open source tools and frameworks. Use Dataproc for data lake modernization, ETL, and secure data science, at scale, integrated with Google Cloud, at a fraction of the cost. https://cloud.google.com/dataproc

upvoted 1 times

omermahgoub 3 years, 8 months ago

To scale the number and size of Apache Spark and Hadoop jobs being run on a local datacenter with the least amount of operations work and code change, you should consider using Google Cloud Dataproc, option B. Google Cloud Dataproc is a fully-managed service that makes it easy to run Apache Spark and Hadoop workloads in the cloud. It is designed to simplify the process of setting up and managing clusters for data processing, and allows you to scale quickly and easily as demand increases. With Cloud Dataproc, you can create and delete clusters in just a few minutes, and you can use the familiar Apache Spark and Hadoop APIs and tools to process data. This means that you can utilize the cloud to scale your workloads with minimal changes to your code and operations work. Option A: Google Cloud Dataflow, option C: Google Compute Engine, and option D: Google Kubernetes Engine, would not be suitable for this use case, as they do not provide the same level of support for running Apache Spark and Hadoop workloads as Cloud Dataproc.

upvoted 2 times

AniketD 3 years, 9 months ago

B. Dataproc is managed Apache Spark and Hadoop in GCP

upvoted 1 times

zr79 3 years, 10 months ago

Dataproc for Hadoop and spark ecosystem

upvoted 1 times

minmin2020 3 years, 10 months ago

B. Google Cloud Dataproc

upvoted 1 times

holerina 3 years, 11 months ago

B data proc for hadoop and spark

upvoted 1 times

Dhiraj03 4 years, 2 months ago

Keyword - Apache Spark and Hadoop jobs - Go with Dataproc

upvoted 1 times

Superr 4 years, 2 months ago

dataproc

upvoted 1 times

Nirca 4 years, 4 months ago

Google Cloud Dataproc == managed Spark and Hadoop service

upvoted 2 times

pakochiu 4 years, 4 months ago

B - Dataproc Lift&Shift of Apache Spark and Hadoop jobs

upvoted 1 times

llanerox 4 years, 6 months ago

B is ok.

upvoted 1 times

Question #19 *Topic 1*

The database administration team has asked you to help them improve the performance of their new database server running on Google Compute Engine. The database is for importing and normalizing their performance statistics and is built with MySQL running on Debian Linux. They have an n1-standard-8 virtual machine with 80 GB of SSD persistent disk.
What should they change to get better performance from this system?

A. Increase the virtual machine's memory to 64 GB

B. Create a new virtual machine running PostgreSQL

C. Dynamically resize the SSD persistent disk to 500 GB

D. Migrate their performance metrics warehouse to BigQuery

E. Modify all of their batch jobs to use bulk inserts into the database

Suggested Answer: *C*
Official Answer: C

Community vote distribution

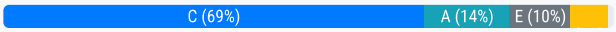

shandy Highly Voted 6 years, 8 months ago
Answer is C because persistent disk performance is based on the total persistent disk capacity attached to an instance and the number of vCPUs that the instance has. Incrementing the persistent disk capacity will increment its throughput and IOPS, which in turn improve the performance of MySQL.
upvoted 73 times

Eroc Highly Voted 6 years, 9 months ago
Assuming that the database is approaching its hardware limits... both options A and C would improve performance, A would increase number of CPUs and memory, but C would increase memory by more. If it a software problem, it is likly it is a hashing problem (the search and sort algorithms are not specific enough to search within the database). This problem would not be fixed just by migrating to PostgreSQL or BigQuery but modifying the inserts would help the situation because it would entail specifications of data lookups. However, it wouldn't help with search performance just inserts and it doesn't help in normalization. So B, D, and E are eliminated. Since statistics is based on sets, the larger the number of sets the better the predictions. This means that the largest amount of memory would not only increase computer performance but also knowledge enhancements. So C beats A.By migration to BigQuery the problem will be fixed for sure. But it requires data migration plan and work. More easier to find root cause of the performance problem and address it.C is ok.When you increase the memory yo need to shutdown the machine, but when you increase the disk, it is not necessary. Answer is B.there isn't "without downtime"I wanted to put letter C.Since its using SQL, so there will be a Maintenance Window, so this change can be implemented during the downtime(also there is no mention that the system should be always avaliable)The IOPS on SSD on GCP increase with the size.
upvoted 37 times

Eroc 6 years, 9 months ago
Assuming that the database is approaching its hardware limits... both options A and C would improve performance, A would increase number of CPUs and memory, but C would increase memory by more. If it a software problem, it is likly it is a hashing problem (the search and sort algorithms are not specific enough to search within the database). This problem would not be fixed just by migrating to PostgreSQL or BigQuery but modifying the inserts would help the situation because it would entail specifications of data lookups. However, it wouldn't help with search performance just inserts and it doesn't help in normalization. So B, D, and E are eliminated. Since statistics is based on sets, the larger the number of sets the better the predictions. This means that the largest amount of memory would not only increase computer performance but also knowledge enhancements. So C beats A.
upvoted 37 times

hitmax87 Most Recent 2 months, 3 weeks ago
By migration to BigQuery the problem will be fixed for sure. But it requires data migration plan and work. More easier to find root cause of the performance problem and address it.
upvoted 1 times

tartar 6 years ago
C is ok.
upvoted 8 times

haroldbenites 4 years, 8 months ago

When you increase the memory yo need to shutdown the machine, but when you increase the disk, it is not necessary. Answer is B.there isn't "without downtime"I wanted to put letter C.Since its using SQL, so there will be a Maintenance Window, so this change can be implemented during the downtime(also there is no mention that the system should be always avaliable)

upvoted 2 times

haroldbenites 4 years, 8 months ago

When you increase the memory yo need to shutdown the machine, but when you increase the disk, it is not necessary. Answer is B.

upvoted 2 times

Dclaiborne41 4 years, 3 months ago

there isn't "without downtime"

upvoted 2 times

haroldbenites 4 years, 8 months ago

I wanted to put letter C.

upvoted 2 times

Mission94 2 years, 2 months ago

Since its using SQL, so there will be a Maintenance Window, so this change can be implemented during the downtime(also there is no mention that the system should be always avaliable)

upvoted 1 times

Matro71 6 years, 1 month ago

The IOPS on SSD on GCP increase with the size.

upvoted 12 times

Johnb9682 4 months, 1 week ago

The question is missing a condition that DBA want "Fast" changes. Otherwise, D makes more sense.

upvoted 1 times

ahmed75.zahran00 7 months ago

https://docs.cloud.google.com/compute/docs/disks/performance The following tables show IOPS and throughput limits for regional Persistent Disk. These conditions apply to the stated limits: Achieving the stated limits requires a compute instance with at least 64 vCPUs and an N2 machine type or an N1 machine type running on the Intel Skylake or newer CPU. Instances using other machine types, such as N2D, have lower performance limits.

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

OLTP D/B performance is depended on IOPs. SSD is the best solution for higher IOPs. In GCP bigger the disk size higher the IOPs.

upvoted 3 times

Hungdv 2 years ago

Choose C

upvoted 1 times

ukivanlamlpi 2 years, 1 month ago

increase size will not increase performance, it either increase RAM or serverless. A or D. if no cost concern will pick D

upvoted 2 times

ashishdwi007 2 years, 7 months ago

I was looking for CloudSQL in options, since it is not there, C is best

upvoted 1 times

hzaoui 2 years, 7 months ago

The fact that the database is used for importing and normalizing performance statistics suggests frequent data insertions. Optimizing this process through bulk inserts directly addresses a likely performance bottleneck.

upvoted 1 times

JohnDohertyDoe 2 years, 7 months ago

Selected Answer: A

The answer according to Google is A. This question is part of the Google's sample questions for the certification.

upvoted 1 times

JohnDohertyDoe 2 years, 7 months ago

The answer according to Google is A. This question is part of the Google's sample questions for the certification.Yes, it is, and says it is C.

upvoted 1 times

ccpmad 2 years, 2 months ago

Yes, it is, and says it is C.

upvoted 1 times

JPA210 2 years, 10 months ago

I see most of the people here replying C, but I do not think that the size of the disk we bring much gains in performance. D, yes, seems to me that will bring much improvements in performance, management, operations and cost. So B, Migrate their performance metrics warehouse to BigQuery

upvoted 4 times

Palan 3 years ago

I would go with option E because Bulk Insert improves performance drastically unless it is been implemented already.

upvoted 1 times

eka_nostra 3 years ago

Increasing disk size will also increase its performance. https://cloud.google.com/compute/docs/disks/performance#optimize_disk_performance

upvoted 4 times

JC0926 3 years, 4 months ago

C. Dynamically resize the SSD persistent disk to 500 GB By increasing the size of the SSD persistent disk, the database server can achieve better performance. A larger SSD persistent disk provides higher IOPS (input/output operations per second) and throughput, allowing for faster read and write operations. This can help improve the performance of the MySQL database server running on the Google Compute Engine instance.

upvoted 3 times

mifrah 3 years, 4 months ago

On another website I found the question with the hint "you are not allowed to reboot the VM before next maintenance window". That makes it more clear --> C.

upvoted 1 times

JC0926 3 years, 5 months ago

E. Modify all of their batch jobs to use bulk inserts into the database: This can be a very effective solution for improving performance. Bulk inserts can greatly reduce the number of round-trips to the database, which can help to minimize latency and improve overall throughput. Therefore, option E is the best choice for improving performance in this scenario.

upvoted 4 times

Jackalski 3 years, 8 months ago

in option C - even increasing disc can gain performance - that will take few months to face new limits. mySQL is not desiged for OLAP/analytics - but OLTP. so I vote on D

upvoted 3 times

Question #20 *Topic 1*

You want to optimize the performance of an accurate, real-time, weather-charting application. The data comes from 50,000 sensors sending 10 readings a second, in the format of a timestamp and sensor reading.
Where should you store the data?

A. Google BigQuery

B. Google Cloud SQL

C. Google Cloud Bigtable

D. Google Cloud Storage

Suggested Answer: *C*
Google Cloud Bigtable is a scalable, fully-managed NoSQL wide-column database that is suitable for both real-time access and analytics workloads. Good for: ✑ Low-latency read/write access ✑ High-throughput analytics ✑ Native time series support Common workloads: ✑ IoT, finance, adtech ✑ Personalization, recommendations ✑ Monitoring ✑ Geospatial datasets ✑ Graphs Reference: https://cloud.google.com/storage-options/

Community vote distribution
C (100%)

victory108 Highly Voted 5 years, 1 month ago
C. Google Cloud Bigtable
upvoted 12 times

victory108 Highly Voted 5 years, 1 month ago
C. Google Cloud BigtableI too got this question in 10-09-22 exam with similar option and result is pass
upvoted 12 times

khadar 3 years, 11 months ago
I too got this question in 10-09-22 exam with similar option and result is pass
upvoted 4 times

FabPanda Most Recent 1 year, 6 months ago
Cloud Bigtable is right solution and correct database choice, which provides high throughput, low latency and scalability for time series data such as this case.
upvoted 2 times

Ekramy_Elnaggar 1 year, 9 months ago
1. High Write Throughput: Bigtable excels at handling high-volume write operations, which is crucial for your application receiving data from 50,000 sensors sending 10 readings per second. 2. Low Latency: Bigtable offers very low latency for read operations, essential for real-time charting and data visualization. 3. Time-Series Data: Bigtable is well-suited for storing and querying time-series data, like your weather sensor readings with timestamps. 4. Scalability: Bigtable can handle massive amounts of data and scale seamlessly as your application grows.
upvoted 3 times

lisabisa 2 years, 5 months ago
Bigtable - NoSQL, high-throughput, low-latency, making it suitable for storing time-series data from sensors
upvoted 2 times

alekonko 3 years, 4 months ago
BigTable is NoSQL for IoT
upvoted 2 times

sivaamum 3 years, 5 months ago
C is correct
upvoted 1 times

omermahgoub 3 years, 8 months ago
To optimize the performance of an accurate, real-time, weather-charting application that receives data from 50,000 sensors sending 10 readings per second, it would be most appropriate to store the data in a distributed, horizontally scalable, NoSQL database such as Google Cloud Bigtable Other options, such as Google BigQuery, Google Cloud SQL, and Google Cloud Storage, may not be as well-suited for handling high volumes of real-time data and may not provide the same level of performance and scalability as Google Cloud Bigtable.

upvoted 3 times

Bry_040706 3 years, 8 months ago

C. Bigtable, IoT data.

upvoted 1 times

AniketD 3 years, 9 months ago

C Bigtable

upvoted 1 times

Mahmoud_E 3 years, 10 months ago

C bigtable right answer

upvoted 1 times

zr79 3 years, 10 months ago

real-time, IoT, time series and huge writes are some of the keywords to look after for Bigtable

upvoted 4 times

minmin2020 3 years, 10 months ago

C. Google Cloud Bigtable

upvoted 1 times

holerina 3 years, 11 months ago

C big table for IOT data

upvoted 1 times

abirroy 3 years, 11 months ago

Google Cloud Bigtable

upvoted 1 times

Dhiraj03 4 years, 2 months ago

Keyword - Timestamp - Big table

upvoted 2 times

szanio 4 years, 2 months ago

Go for c

upvoted 1 times

Nirca 4 years, 4 months ago

C. Google Cloud Bigtable is the Best Practice option

upvoted 1 times

Question #21 *Topic 1*

Your company's user-feedback portal comprises a standard LAMP stack replicated across two zones. It is deployed in the us-central1 region and uses autoscaled managed instance groups on all layers, except the database. Currently, only a small group of select customers have access to the portal. The portal meets a
99,99% availability SLA under these conditions. However next quarter, your company will be making the portal available to all users, including unauthenticated users. You need to develop a resiliency testing strategy to ensure the system maintains the SLA once they introduce additional user load.
What should you do?

A. Capture existing users input, and replay captured user load until autoscale is triggered on all layers. At the same time, terminate all resources in one of the zones

B. Create synthetic random user input, replay synthetic load until autoscale logic is triggered on at least one layer, and introduce ג€chaosג€ to the system by terminating random resources on both zones

C. Expose the new system to a larger group of users, and increase group size each day until autoscale logic is triggered on all layers. At the same time, terminate random resources on both zones

D. Capture existing users input, and replay captured user load until resource utilization crosses 80%. Also, derive estimated number of users based on existing user's usage of the app, and deploy enough resources to handle 200% of expected load

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (74%) A (19%) (4 (4

jcmoranp Highly Voted 6 years, 9 months ago
resilience test is not about load, is about terminate resources and service not affected. Think it's B. The best for resilience in to introduce chaos in the infraestructure
upvoted 95 times

jcmoranp Highly Voted 6 years, 9 months ago
resilience test is not about load, is about terminate resources and service not affected. Think it's B. The best for resilience in to introduce chaos in the infraestructureIsn't A superior in one way. It will demonstrate that the app is regionally redundant by demonstrating it can survive the loss of an entire zone. B only demonstrates the app is zonally redundant and can lose a random instance here and there within individual zones which is not that resilient. Thoughts?No. It is only terminating the service in ONE zone. B caters for terminating the service in both zones randomly. You want to be able to test resiliency when either zone has an outage.I agree with @jcmoranp, B) is correct for more info - https://cloud.google.com/solutions/scalable-and-resilient-apps#test_your_resilience
upvoted 95 times

AWSPro24 4 years, 9 months ago
Isn't A superior in one way. It will demonstrate that the app is regionally redundant by demonstrating it can survive the loss of an entire zone. B only demonstrates the app is zonally redundant and can lose a random instance here and there within individual zones which is not that resilient. Thoughts?
upvoted 9 times

AWSPro24 4 years, 9 months ago
Isn't A superior in one way. It will demonstrate that the app is regionally redundant by demonstrating it can survive the loss of an entire zone. B only demonstrates the app is zonally redundant and can lose a random instance here and there within individual zones which is not that resilient. Thoughts?No. It is only terminating the service in ONE zone. B caters for terminating the service in both zones randomly. You want to be able to test resiliency when either zone has an outage.
upvoted 9 times

0xE8D4A51000 3 years, 10 months ago
No. It is only terminating the service in ONE zone. B caters for terminating the service in both zones randomly. You want to be able to test resiliency when either zone has an outage.
upvoted 7 times

rockstar9622 6 years, 7 months ago

I agree with @jcmoranp, B) is correct for more info - https://cloud.google.com/solutions/scalable-and-resilient-apps#test_your_resilience

upvoted 20 times

OSNG 5 years, 9 months ago

Will go with A. Reason: 1. SLA in question is about the Availability (The portal meets a 99,99% availability SLA under these conditions.) therefore maintaining SLA means Availability. 2. Its a user-feedback portal and type of user input is going to be similar or same (A is capturing the user input and replaying it). Why not B: The infrastructure is using MIG (Instances created using templates) most likely to be used with Health Check and killing random VMs cannot test the availability (neither affect the availability as health check will immediately kill the effected Instances and create the other one.) Why not D: SLA is about Availability not reliability or scaling. (As all of it does work hand to hand but still major focus should be on availability.) --- IF AGREE PLEASE UP VOTE TO MAKE IT CLEAR FOR THE OTHERS --- Thank you.valuable input in terms of 'availability'. did you select this answer in exam too?We are talking about resilience testing where as SLA is an argument of the system.And resilience means the capacity to recover from failure.A ensures the app can withstand the loss of a whole Zone which I think is important as well.also, it's mentionned that there is only one DB, if randomly we kill the DB, all the system will be KO

upvoted 63 times

OSNG 5 years, 9 months ago

Will go with A. Reason: 1. SLA in question is about the Availability (The portal meets a 99,99% availability SLA under these conditions.) therefore maintaining SLA means Availability. 2. Its a user-feedback portal and type of user input is going to be similar or same (A is capturing the user input and replaying it).

Why not B: The infrastructure is using MIG (Instances created using templates) most likely to be used with Health Check and killing random VMs cannot test the availability (neither affect the availability as health check will immediately kill the effected Instances and create the other one.) Why not D: SLA is about Availability not reliability or scaling. (As all of it does work hand to hand but still major focus should be on availability.)

--- IF AGREE PLEASE UP VOTE TO MAKE IT CLEAR FOR THE OTHERS --- Thank you.

upvoted 63 times

bolu 5 years, 7 months ago

valuable input in terms of 'availability'. did you select this answer in exam too?

upvoted 1 times

amxexam 4 years, 11 months ago

We are talking about resilience testing where as SLA is an argument of the system.

upvoted 1 times

amxexam 4 years, 11 months ago

We are talking about resilience testing where as SLA is an argument of the system.And resilience means the capacity to recover from failure.

upvoted 1 times

amxexam 4 years, 11 months ago

And resilience means the capacity to recover from failure.

upvoted 1 times

AWSPro24 4 years, 9 months ago

A ensures the app can withstand the loss of a whole Zone which I think is important as well.

upvoted 1 times

bargou 2 years, 6 months ago

also, it's mentionned that there is only one DB, if randomly we kill the DB, all the system will be KO

upvoted 1 times

41168b9 Most Recent 8 months, 2 weeks ago

Option A combines the two most critical types of testing required: 1. Load Testing (Scaling and Performance) Action: "Capture existing users input, and replay captured user load until autoscale is triggered on all layers." 2. Chaos Engineering / Resiliency Testing (Zone Failure) Action: "At the same time, terminate all resources in one of the zones."

upvoted 1 times

scialappa 1 year, 4 months ago

I'm a bit puzzled by none selecting C. Why shouldn't C be the best options. Aren't ABD only partially fulfilling the requirements, whilst C marks them all? Real users to test (most realistic), allows stress testing by increasing load gradually, randomly shuts resources so includes Cahos Eng. Any feedback here please?

upvoted 2 times

david_tay 1 year, 5 months ago

Answer is likely A as Gemini said so, and some others here agree on the logic.

upvoted 1 times

FabPanda 1 year, 6 months ago

To perform resilience test and auto scale testing B is correct choice.

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Synthetic Load Generation: Creating synthetic user input allows you to simulate a wide range of user behavior and load patterns, including spikes and sustained high traffic. This helps you test the system's ability to scale and handle unexpected loads. 2. Autoscaling Validation: By replaying the synthetic load, you can verify that the autoscaling logic is working correctly across all layers of the LAMP stack. This ensures that the system can dynamically adjust resources to meet demand. 3. Chaos Engineering: Introducing chaos by terminating random resources simulates real-world failures and helps you test the system's resilience to unexpected disruptions. This is crucial for maintaining the 99.99% availability SLA. 4. Controlled Environment: This approach allows you to conduct testing in a controlled environment without impacting real users. You can gradually increase the load and introduce chaos in a measured way to identify weaknesses and improve resilience.

upvoted 3 times

potorange 1 year, 11 months ago

A: requirements states "all layers" and "resiliency testing"

upvoted 1 times

Robert0 2 years, 2 months ago

I would go with A. This solution test autoscale policy of each layer (not only one as option B refers). Also, it propose a regional shutdown. This is a very good test commonly requested if your application is geo-redundant. In crontast, option B propose random termination of resources, not a bad practice but a little bit vague that can be implemented terrible wrong (for example you do not kill the interesting services or you kill the same service in both regions, thus generating a blackout)

upvoted 1 times

lisabisa 2 years, 5 months ago

We need to do 1. load testing and 2. reliability test (failover redundency) B does both A only tests one zone C impacts real user experience D 200% not necessary

upvoted 4 times

AdityaGupta 2 years, 10 months ago

You need to develop a resiliency testing strategy to ensure the system maintains the SLA once they introduce additional user load. Need to maintain SLA of 99.9% means multiple zones, resilience means fault tolerance. Teminating all resources in one zone is also creating a chaos.

upvoted 1 times

heretolearnazure 2 years, 12 months ago

B is correct

upvoted 1 times

VaraSrinvas 3 years, 2 months ago

Option D is the best resiliency testing strategy in this scenario as it ensures that the system is tested with actual user data, takes into account the expected increase in user load, and ensures that the system is adequately scaled to handle the anticipated load.Do not agree. B is 100% correctBut this would assume that the user load will not change; plus, the current application is visible only to a small group of select customers - this is the current production setup. The deployment should be prepared for all existing users plus unauthenticated users, and the load increase is unknown, so testing for 200% of "expected load" is very ambiguous.

upvoted 2 times

VaraSrinvas 3 years, 2 months ago

Selected Answer: D

Option D is the best resiliency testing strategy in this scenario as it ensures that the system is tested with actual user data, takes into account the expected increase in user load, and ensures that the system is adequately scaled to handle the anticipated load.

upvoted 2 times

didek1986 3 years ago

Do not agree. B is 100% correct

upvoted 1 times

jrisl1991 2 years, 10 months ago

But this would assume that the user load will not change; plus, the current application is visible only to a small group of select customers - this is the current production setup. The deployment should be prepared for all existing users plus unauthenticated users, and the load increase is unknown, so testing for 200% of "expected load" is very ambiguous.

upvoted 1 times

JC0926 3 years, 4 months ago

B. Create synthetic random user input, replay synthetic load until autoscale logic is triggered on at least one layer, and introduce ɹ€chaosɹ€ to the system by terminating random resources on both zones. By creating synthetic random user input and replaying the load, you can simulate the

expected increased user traffic and trigger the autoscale logic on different layers of the application. Introducing chaos to the system by terminating random resources in both zones helps test the resiliency and redundancy of the system under stress. This strategy will help ensure that the system can maintain the 99.99% availability SLA when subjected to additional user load.

upvoted 5 times

telp 3 years, 4 months ago

chaos == test resilience for google

upvoted 2 times

Deb2293 3 years, 5 months ago

This is chaos engineering used by Netflix. https://netflixtechblog.com/tagged/chaos-engineering

upvoted 2 times

roaming_panda 3 years, 7 months ago

chaos == checking resilience

upvoted 4 times

Question #22 *Topic 1*

One of the developers on your team deployed their application in Google Container Engine with the Dockerfile below. They report that their application deployments are taking too long.
You want to optimize this Dockerfile for faster deployment times without adversely affecting the app's functionality.
Which two actions should you take? (Choose two.)

```
FROM ubuntu:16.04
COPY . /src
RUN apt-get update && apt-get install -y python python-pip
RUN pip install -r requirements.txt
```

A. Remove Python after running pip

B. Remove dependencies from requirements.txt

C. Use a slimmed-down base image like Alpine Linux

D. Use larger machine types for your Google Container Engine node pools

E. Copy the source after he package dependencies (Python and pip) are installed

Suggested Answer: *C, E*
The speed of deployment can be changed by limiting the size of the uploaded app, limiting the complexity of the build necessary in the Dockerfile, if present, and by ensuring a fast and reliable internet connection. Note: Alpine Linux is built around musl libc and busybox. This makes it smaller and more resource efficient than traditional GNU/Linux distributions. A container requires no more than 8 MB and a minimal installation to disk requires around 130 MB of storage. Not only do you get a fully-fledged Linux environment but a large selection of packages from the repository. Reference: https://groups.google.com/forum/#!topic/google-appengine/hZMEkmmObDU
https://www.alpinelinux.org/about/

Community vote distribution



aviratna Highly Voted 5 years, 1 month ago
C & E: C: Smaller the base image with minimum dependency faster the container will start E: Docker image build uses caching. Docker Instructions sequence matter because application's dependencies change less frequently than the Python code which will help to reuse the cached layer of dependency and only add new layer for code change for Python Source code.
upvoted 62 times

vincy2202 Highly Voted 4 years, 9 months ago
C & E are the correct answers. Kindly refer - https://www.docker.com/blog/intro-guide-to-dockerfile-best-practices/
upvoted 12 times

MrAZ105 Most Recent 1 year, 2 months ago
CE are correct
upvoted 1 times

halifax 1 year, 5 months ago
A. Remove Python after running pip: This translates to using a multi-stage build. C. Use a slimmed-down base image like Alpine Linux: While the multi-stage build (option A) is the most powerful optimization, you can also change the final stage by using Alpine. A smaller base image means less data to download and store, resulting in faster builds and deployments.
upvoted 1 times

zacrparsons 1 year, 7 months ago
do go against some others here i think B makes more sense that C C wont actually decrease build time as the docker repo should already have the base image downloaded. where as if there are indeed dependancies in requirements.txt that can be removed then B would reduce the build time for the install.
upvoted 2 times

Ekramy_Elnaggar 1 year, 9 months ago

C. Use a slimmed-down base image like Alpine Linux: The ubuntu:16.04 image is a full-fledged operating system, which means it's larger and takes longer to download and build. Alpine Linux is a minimal distribution designed for containers, resulting in significantly smaller images and faster deployments. E. Copy the source after the package dependencies (Python and pip) are installed: Docker builds images in layers. Each RUN, COPY, and ADD instruction creates a new layer. By copying the source code after installing dependencies, you can take advantage of Docker's caching mechanism. If your source code changes, only the layers related to the source code need to be rebuilt, not the layers related to dependencies.

upvoted 4 times

alessandroGPC 2 years, 3 months ago

Hi everyone, I have a doubt. The question talks about "deployment" not "build". CE are more correct, in my opinion, to accelerate the build phase (and application management) rather than the simple deployment (with docker swarm, simple docker, kubernetes etc etc)

upvoted 1 times

nickcin77 2 years, 8 months ago

C&E C - Use small images. E - "Try to make expensive steps appear near the beginning of the Dockerfile. Steps that change often should appear near the end of the Dockerfile" https://docs.docker.com/build/cache/

upvoted 3 times

jrisl1991 2 years, 10 months ago

C is an obvious choice, as an optimized image will be much better for a container than the one in the file. For E, I found an explanation here about COPY that helped me confirm that's the other solution - https://test-dockerrr.readthedocs.io/en/latest/userguide/eng-image/dockerfile_best-practices/#:~:text=required%20files%20change.-,For%20example%3A,-COPY%20requirements.txt.

upvoted 1 times

poplot321 3 years, 1 month ago

I know C & E makes the most sense, but the Dockerfile is has a bug. If you don't copy requirements.txt from source, you have nothing to pip install. The line after FROM should be: COPY requirements.txt requirements.txt

upvoted 4 times

Badri9898 3 years, 4 months ago

The two actions that should be taken to optimize the Dockerfile for faster deployment times without adversely affecting the app's functionality are:

B. Remove dependencies from requirements.txt: The requirements.txt file should only contain necessary dependencies to reduce the number of packages to be installed.

C. Use a slimmed-down base image like Alpine Linux: The Alpine Linux image is smaller than Ubuntu and has a smaller attack surface, which reduces the container's build time and image size.

Therefore, options A, D, and E are not correct as they do not directly address the issue of slow deployment times caused by a bloated Dockerfile.

upvoted 1 times

Badri9898 3 years, 4 months ago

The two actions that should be taken to optimize the Dockerfile for faster deployment times without adversely affecting the app's functionality are: B. Remove dependencies from requirements.txt: The requirements.txt file should only contain necessary dependencies to reduce the number of packages to be installed. C. Use a slimmed-down base image like Alpine Linux: The Alpine Linux image is smaller than Ubuntu and has a smaller attack surface, which reduces the container's build time and image size. Therefore, options A, D, and E are not correct as they do not directly address the issue of slow deployment times caused by a bloated Dockerfile.But how do you know that requirements.txt has other dependencies? We don't have the file to confirm if there are unnecessary dependencies. Assuming that we only have the necessary ones, removing dependencies would break the deployment.

upvoted 1 times

jrisl1991 2 years, 10 months ago

But how do you know that requirements.txt has other dependencies? We don't have the file to confirm if there are unnecessary dependencies. Assuming that we only have the necessary ones, removing dependencies would break the deployment.

upvoted 1 times

alekonko 3 years, 4 months ago

C: use smaller image decrease pull time E: optimize build time using previous cache layer image. generate new layer only for a different app code and requirements A: can't remove python B: the developer choose right deps D: changing istance type don't directly reduce deploy time

upvoted 2 times

omermahgoub 3 years, 8 months ago

C & E Using a slimmed-down base image like Alpine Linux can help reduce the size of your Docker image, which can lead to faster deployment times. Alpine Linux is a lightweight Linux distribution that is often used as a base image for Docker images because of its small size.

Additionally, copying the source code after installing the package dependencies can help reduce the image build time because the dependencies will only need to be installed once, rather than every time the source code is changed. This can lead to faster deployment times because the image build process will be faster.

upvoted 7 times

omermahgoub 3 years, 8 months ago

C & E Using a slimmed-down base image like Alpine Linux can help reduce the size of your Docker image, which can lead to faster deployment times. Alpine Linux is a lightweight Linux distribution that is often used as a base image for Docker images because of its small size. Additionally, copying the source code after installing the package dependencies can help reduce the image build time because the dependencies will only need to be installed once, rather than every time the source code is changed. This can lead to faster deployment times because the image build process will be faster.It is not recommended to remove dependencies from the requirements.txt file or remove Python after running pip, as this could adversely affect the functionality of the application. Similarly, using larger machine types for your Google Container Engine node pools may not directly affect the deployment times of your application, as the deployment times are primarily dependent on the size and complexity of the Docker image being deployed.

upvoted 7 times

omermahgoub 3 years, 8 months ago

It is not recommended to remove dependencies from the requirements.txt file or remove Python after running pip, as this could adversely affect the functionality of the application. Similarly, using larger machine types for your Google Container Engine node pools may not directly affect the deployment times of your application, as the deployment times are primarily dependent on the size and complexity of the Docker image being deployed.

upvoted 1 times

sfsdeniso 3 years, 10 months ago

B & E base image is already cached - so no improvement in build time B is about removing unnecessary dependencies and not about all of them i remember saw this question on google's site - BE are correct

upvoted 2 times

AzureDP900 3 years, 10 months ago

CE and is perfect

upvoted 1 times

minmin2020 3 years, 10 months ago

C. Use a slimmed-down base image like Alpine Linux E. Copy the source after he package dependencies (Python and pip) are installed

upvoted 1 times

backhand 4 years ago

vote C, E https://cloud.google.com/architecture/best-practices-for-building-containers

upvoted 1 times

Question #23 *Topic 1*

Your solution is producing performance bugs in production that you did not see in staging and test environments. You want to adjust your test and deployment procedures to avoid this problem in the future.
What should you do?

A. Deploy fewer changes to production

B. Deploy smaller changes to production

C. Increase the load on your test and staging environments

D. Deploy changes to a small subset of users before rolling out to production

Suggested Answer: *D*
Official Answer: D

Community vote distribution
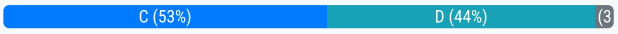

ghitesh Highly Voted 6 years, 7 months ago
Question Statement: You want to adjust your test and deployment procedures to avoid this problem in the future

So based on this, I think the option "C" is correct, since it is the only one talking about doing changes in the test environment.
upvoted 88 times

ghitesh Highly Voted 6 years, 7 months ago
Question Statement: You want to adjust your test and deployment procedures to avoid this problem in the future So based on this, I think the option "C" is correct, since it is the only one talking about doing changes in the test environment.The answer is D"Your solution is producing performance bugs in production..." - I don't see how "D" would help to detect performance bugs. - "C" looks more adequate.C. Increase the load on your test and staging environments. As you have pointed out in "Question Statement", I do not see C covering "deployment procedures". Test and Staging environment is more on testing, but not about deployment procedure to production. So, the only option that cover test and deployment is D. (Yes, kind of unacceptable to have the users to do "testing", but we make it "ok" by calling it "canary deployment")With canary deployment we expose the new version to a small portion of users. With this approach maybe we don't see performance bugs in the canary release, since we don't have the 100% of traffic on the canary. But when we migrate the 100% of traffic to the new release (previous canary) we can see performance bugs.There is no indication given anywhere that the load is the problem or that the bugs are a result of load and not some other issue encountered when using a specific feature."Performance" bug is a result of Load
upvoted 88 times

Urban_Life 4 years, 8 months ago
The answer is D
upvoted 9 times

RegisFTM 4 years, 7 months ago
"Your solution is producing performance bugs in production..." - I don't see how "D" would help to detect performance bugs. - "C" looks more adequate.
upvoted 19 times

VedaSW 5 years, 10 months ago
C. Increase the load on your test and staging environments.

As you have pointed out in "Question Statement", I do not see C covering "deployment procedures". Test and Staging environment is more on testing, but not about deployment procedure to production.

So, the only option that cover test and deployment is D. (Yes, kind of unacceptable to have the users to do "testing", but we make it "ok" by calling it "canary deployment")
upvoted 20 times

VedaSW 5 years, 10 months ago
C. Increase the load on your test and staging environments. As you have pointed out in "Question Statement", I do not see C covering "deployment procedures". Test and Staging environment is more on testing, but not about deployment procedure to production. So, the only option that cover test

and deployment is D. (Yes, kind of unacceptable to have the users to do "testing", but we make it "ok" by calling it "canary deployment")With canary deployment we expose the new version to a small portion of users. With this approach maybe we don't see performance bugs in the canary release, since we don't have the 100% of traffic on the canary. But when we migrate the 100% of traffic to the new release (previous canary) we can see performance bugs.

upvoted 20 times

francescogugliottagm 2 years, 10 months ago

With canary deployment we expose the new version to a small portion of users. With this approach maybe we don't see performance bugs in the canary release, since we don't have the 100% of traffic on the canary. But when we migrate the 100% of traffic to the new release (previous canary) we can see performance bugs.

upvoted 10 times

Sephethus 2 years, 2 months ago

There is no indication given anywhere that the load is the problem or that the bugs are a result of load and not some other issue encountered when using a specific feature."Performance" bug is a result of Load

upvoted 3 times

Sephethus 2 years, 2 months ago

There is no indication given anywhere that the load is the problem or that the bugs are a result of load and not some other issue encountered when using a specific feature.

upvoted 3 times

alihabib 1 year, 8 months ago

"Performance" bug is a result of Load

upvoted 3 times

Eroc 6 years, 9 months ago

A wouldn't prevent the bugs, it would just avoid them. B would help with root-cause analysis because it'd be a smaller change to review. C would test the performance of the system at its peak processing rates, so this assumes the bugs in production only occur because of usage. D would allow you to test the new code against smaller user sets to see if it occurs then, and if it still does you know it is not because of more user responses. So it's a tossup between C and D, D would be the cheaper/quicker answer so I'd choose D first then C if it's because of usage.

upvoted 39 times

Eroc 6 years, 9 months ago

A wouldn't prevent the bugs, it would just avoid them. B would help with root-cause analysis because it'd be a smaller change to review. C would test the performance of the system at its peak processing rates, so this assumes the bugs in production only occur because of usage. D would allow you to test the new code against smaller user sets to see if it occurs then, and if it still does you know it is not because of more user responses. So it's a tossup between C and D, D would be the cheaper/quicker answer so I'd choose D first then C if it's because of usage.D, canary rolloutIt has nothing to do with the "performance bugs"According to the question, [Your solution is producing "performance" bugs in production], so I think it is about the load. Plus canary test will not reproduce the bugs related to high load, I vote for CThe question is about the performance of the existing Code that they did not detect in Test environments . This is not about new API release . In order to test the performance they should increase the load in test environment and hence answer C.C is the best

upvoted 39 times

nitinz 5 years, 5 months ago

D, canary rollout

upvoted 7 times

nitinz 5 years, 5 months ago

D, canary rolloutIt has nothing to do with the "performance bugs"

upvoted 7 times

michael_m 4 years ago

It has nothing to do with the "performance bugs"

upvoted 1 times

michael_m 4 years ago

According to the question, [Your solution is producing "performance" bugs in production], so I think it is about the load. Plus canary test will not reproduce the bugs related to high load, I vote for C

upvoted 2 times

Sreekey 6 years ago

The question is about the performance of the existing Code that they did not detect in Test environments . This is not about new API release . In order to test the performance they should increase the load in test environment and hence answer C.

upvoted 18 times

AzureDP900 3 years, 10 months ago

C is the best

upvoted 2 times

tim777 Most Recent 2 months, 2 weeks ago

D because Canary Deployment, Not C because load testing in staging does not capture all the complexity of the production environment.

upvoted 1 times

WayneHsu 3 months, 3 weeks ago

D is the answer.

upvoted 1 times

Wayne90040 3 months, 3 weeks ago

D is the answer

upvoted 1 times

Johnb9682 4 months, 1 week ago

Should duplicate the production like env in staging

upvoted 1 times

CyberPeter 6 months, 2 weeks ago

Question ask to adjust test environment, not to fix production bugs, so C

upvoted 1 times

Dusan93 7 months ago

D is correct

upvoted 1 times

mahi_h 1 year, 2 months ago

I am going with the key words in the question "adjust your test and deployment procedures", Increasing the load on test/stage env does not realistically adjust any test procedure. Besides that, in my experience as Infra SME, we never design capacity of non-prod environment same as prod (always 50% or even 25%), so amplifying load will not provide "production" like behavior. Option D sounds like making "adjustments" in your procedure.

upvoted 2 times

khadar2001 1 year, 4 months ago

confirmed from practice Exam

upvoted 2 times

Yass92 1 year, 5 months ago

performance = load

upvoted 1 times

halifax 1 year, 5 months ago

No matter how efficiently you load tests in the development stage, you will never catch 100% of possible bugs. This is why CI/CD was invented (option D: continuous development).

upvoted 2 times

hpf97 1 year, 6 months ago

I have that king of questions as C and D would be correct. But as it is said to "adjust", I would prefer C than setting a new canary deployment strategy. Furthermore in case D would a small set of users rises up the performance issue ? not sure.

upvoted 1 times

RVivek 1 year, 7 months ago

Performance Bug (not functional). Hence increasing the load on test environment will help to identify and ifx the bugs

upvoted 1 times

deep316 1 year, 8 months ago

Bugs are related to performance. So you would need to perform performance testing thoroughly.

upvoted 1 times

Daniaw 1 year, 8 months ago

C could be correct but this can help identify performance bottlenecks, but it might not fully replicate the complexity and unpredictability of real-world production traffic. So D is the most correct one.

upvoted 2 times

alihabib 1 year, 8 months ago

"Performance" bug, not a "Development" Bug, so it more aligns with C, to increase the load to determine performance related bugs

upvoted 1 times

Question #24 *Topic 1*

A small number of API requests to your microservices-based application take a very long time. You know that each request to the API can traverse many services.
You want to know which service takes the longest in those cases.
What should you do?

A. Set timeouts on your application so that you can fail requests faster

B. Send custom metrics for each of your requests to Stackdriver Monitoring

C. Use Stackdriver Monitoring to look for insights that show when your API latencies are high

D. Instrument your application with Stackdriver Trace in order to break down the request latencies at each microservice

Suggested Answer: *D*
Reference: https://cloud.google.com/trace/docs/quickstart#find_a_trace

Community vote distribution
D (100%)

euclid Highly Voted 6 years, 7 months ago
D is correct !
upvoted 23 times

euclid Highly Voted 6 years, 7 months ago
D is correct !D is okD, trace is just for latency testing.
upvoted 23 times

tartar 6 years ago
D is ok
upvoted 9 times

nitinz 5 years, 5 months ago
D, trace is just for latency testing.
upvoted 4 times

omermahgoub 3 years, 8 months ago
D. Instrument your application with Stackdriver Trace in order to break down the request latencies at each microservice Stackdriver Trace is a distributed tracing system that allows you to understand the relationships between requests and the various microservices that they touch as they pass through your application. By instrumenting your application with Stackdriver Trace, you can get a detailed breakdown of the latencies at each microservice, which can help you identify which service is taking the longest in those cases where a small number of API requests take a very long time. Setting timeouts on your application or sending custom metrics to Stackdriver Monitoring may not provide the level of detail that you need to identify the specific service that is causing the latency issues. Looking for insights in Stackdriver Monitoring may also not provide the necessary level of detail, as it may not show the individual latencies at each microservice.
upvoted 11 times

hpf97 Most Recent 1 year, 6 months ago
To identify latency in a microservice chain, only Trace can be useful, so D
upvoted 1 times

Ishu_awsguy 1 year, 7 months ago
The service is named Cloud trace now
upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago
1. Distributed Tracing: Stackdriver Trace (now called Cloud Trace) is specifically designed to analyze the performance of requests as they travel through multiple services in a microservices architecture. It helps you pinpoint exactly where latency is occurring. 2. Detailed Breakdown: Cloud Trace will provide a detailed breakdown of how long each microservice takes to process the request, allowing you to identify the bottleneck quickly. 3. Visualization: Cloud Trace provides visualizations like flame graphs and waterfall diagrams that make it easy to see the flow of the request and identify slow services.
upvoted 3 times

Hungdv 2 years ago

Choose D

upvoted 1 times

Robert0 2 years, 2 months ago

This question is not updated. It will be refered as Cloud Trace as part of Google Cloud Operation suite

upvoted 4 times

Robert0 2 years, 2 months ago

This question is not updated. Stackdriver is now called "Google Cloud operations"

upvoted 2 times

eka_nostra 3 years ago

Trace is a signal that can be used to follow the program's data and flow, including the duration of the program's components.

upvoted 2 times

LaxmanTiwari 3 years, 2 months ago

D should be correct, the headline in GC trace documentation says it all: "Cloud Trace is a distributed tracing system for Google Cloud that collects latency data from applications and displays it in near real-time in the Google Cloud Console."

upvoted 2 times

LaxmanTiwari 3 years, 2 months ago

D should be correct, the headline in GC trace documentation says it all: "Cloud Trace is a distributed tracing system for Google Cloud that collects latency data from applications and displays it in near real-time in the Google Cloud Console."even got confused with the C ... after reading the conversation and reference doc D make sense to me.

upvoted 2 times

LaxmanTiwari 3 years, 2 months ago

even got confused with the C ... after reading the conversation and reference doc D make sense to me.

upvoted 1 times

alekonko 3 years, 4 months ago

D is the correct answer

upvoted 1 times

MestreCholas 3 years, 5 months ago

Why not C?

upvoted 1 times

surajkrishnamurthy 3 years, 8 months ago

D is the correct answer

upvoted 1 times

AniketD 3 years, 9 months ago

Stacdrive Trace would trace the APIs and helps to identify the bottleneck

upvoted 1 times

kchandank 3 years, 9 months ago

D is correct trace would report to the latency

upvoted 1 times

Mahmoud_E 3 years, 10 months ago

D is correct trace would report to the latency

upvoted 1 times

minmin2020 3 years, 10 months ago

Selected Answer: D

D. Instrument your application with Stackdriver Trace in order to break down the request latencies at each microservice

upvoted 2 times

minmin2020 3 years, 10 months ago

D. Instrument your application with Stackdriver Trace in order to break down the request latencies at each microservicethis is the best option to find more details

upvoted 2 times

AzureDP900 3 years, 10 months ago

this is the best option to find more details

upvoted 1 times

Question #25 *Topic 1*

During a high traffic portion of the day, one of your relational databases crashes, but the replica is never promoted to a master. You want to avoid this in the future.
What should you do?

A. Use a different database

B. Choose larger instances for your database

C. Create snapshots of your database more regularly

D. Implement routinely scheduled failovers of your databases

Suggested Answer: *D*
Official Answer: D

Community vote distribution
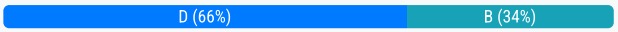

Narigdo Highly Voted 6 years, 8 months ago
Answer is DYep, +1 for D
upvoted 103 times

Narigdo Highly Voted 6 years, 8 months ago
Answer is D
upvoted 103 times

Jos 6 years, 8 months ago
Yep, +1 for D
upvoted 20 times

Eroc 6 years, 9 months ago
@chiar, I agree the question i s not clear. In GCP larger instances have larger number of CPUs, Memory and come with their own private network. So increases the instance size would help prevent the need for failover during high traffic times. However, routinely scheduled failovers would allow the team to test the failover when it is not requried. This would make sure it is working when it is required.
upvoted 45 times

Eroc 6 years, 9 months ago
@chiar, I agree the question i s not clear. In GCP larger instances have larger number of CPUs, Memory and come with their own private network. So increases the instance size would help prevent the need for failover during high traffic times. However, routinely scheduled failovers would allow the team to test the failover when it is not requried. This would make sure it is working when it is required.exactly how do you know the optimal size. it will be a guess. answer should be D
upvoted 45 times

Shariq 6 years, 8 months ago
exactly how do you know the optimal size. it will be a guess. answer should be D
upvoted 9 times

Johnb9682 Most Recent 4 months, 1 week ago
Should always routate
upvoted 1 times

mahi_h 1 year, 2 months ago
I'll go with option D based on real life experience. As part of managing environments, failover test for the database is part of regular Infra activity (at least once a year)
upvoted 1 times

e5966f2 1 year, 2 months ago
Not D because replica is also crashing (why else would it not be promoted). We don't build a layer to fix an underlying layer that isn't working. We fix why the layer isn't working so the answer is B.
upvoted 2 times

khadar2001 1 year, 4 months ago
Here the issue is existing instance doesn't have enough resources. Even though we ensured that the replica is promoted it too will fail.

upvoted 1 times

hpf97 1 year, 6 months ago

D : Chaos testing

upvoted 1 times

alihabib 1 year, 6 months ago

D - I think question is related to resiliency & disaster recovery, no matter how big or small the instance is.. Is the application able to achieve its RTO & RPO ?.. We don't know what cause the crash, but we could have avoided it if we had a Routine Scheduled Failovers Check done in past rather than anticipating & waiting for a real crash.

upvoted 1 times

alihabib 1 year, 8 months ago

Only a routine check, would indicate a best practice. Larger Instance doesn't guarantee a failover would be prevented

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Proactive Testing: Regularly scheduled failovers proactively test your database's ability to recover and ensure the replica can successfully take over as the master. This helps identify and address any issues in the failover process before a real crisis occurs. 2. Confidence in Failover: By routinely testing failover, you gain confidence that your system can recover automatically in case of a primary database failure, minimizing downtime and data loss. 3. Improved Recovery Time: Regular failovers help optimize the recovery process, reducing the time it takes to switch to the replica.

upvoted 3 times

beagle_Masato 1 year, 9 months ago

Answer is D

upvoted 1 times

james2033 2 years, 2 months ago

- **A. Use a different database**: Simply switching to a different database does not inherently solve the problem of failover and promotion mechanisms. The issue is more about the setup and management of failover strategies rather than the specific database technology used. - **B. Choose larger instances for your database**: While using larger instances might improve performance and potentially reduce the risk of a crash due to resource constraints, it does not address the failover mechanism. The key problem is the replica not being promoted, which larger instances alone won't fix. - **C. Create snapshots of your database more regularly**: Regular snapshots are useful for backups and recovery but do not help with automatic failover and high availability. Snapshots do not ensure that a replica will be promoted to a master if the primary fails.

upvoted 3 times

huuthanhdlv 2 years, 2 months ago

I think the answer is B, as the most important thing is customer experience. We can NOT expect database fails as a normal event which have direct customer and business impacts (if current data base fail because of load then replica database may fail as well). Of course we need to setup the failover process to work, but the more important task will be to increase database load first, thus I choose B.

upvoted 1 times

Jen3 2 years, 5 months ago

I think the answer is D because we can not assume the crash was due to load.

upvoted 1 times

Jen3 2 years, 5 months ago

I think the answer is D because we can not assume the crash was due to load.Further the issue isn't that a crash occurred, the issue is the contingency that was in place didn't kick in. Hence D as my choice.

upvoted 1 times

Jen3 2 years, 5 months ago

Further the issue isn't that a crash occurred, the issue is the contingency that was in place didn't kick in. Hence D as my choice.

upvoted 2 times

ashishdwi007 2 years, 6 months ago

D is correct with given choice, because what if larger instance size is also failed. How ever question ignores logging and networking completely, the best way is to use logging why the DB is crashed, failover is just a remedy to solve the issues at current, not solving the problem itself.

upvoted 2 times

discuss24 2 years, 7 months ago

It is important to identify key words, the issue is replica not being promoted when primary instance fails. Regular testing would identify issue

upvoted 2 times

AWS_Sam 2 years, 7 months ago

The correct answer is D The question is asking how to avoid the situation of not failing over. The answer is to test the failover procedure. The question does *NOT*ask about what happens in the future and whether the secondary node is large enough.

upvoted 1 times

Question #26 *Topic 1*

Your organization requires that metrics from all applications be retained for 5 years for future analysis in possible legal proceedings.
Which approach should you use?

A. Grant the security team access to the logs in each Project

B. Configure Stackdriver Monitoring for all Projects, and export to BigQuery

C. Configure Stackdriver Monitoring for all Projects with the default retention policies

D. Configure Stackdriver Monitoring for all Projects, and export to Google Cloud Storage

Suggested Answer: *B*
Stackdriver Logging provides you with the ability to filter, search, and view logs from your cloud and open source application services. Allows you to define metrics based on log contents that are incorporated into dashboards and alerts. Enables you to export logs to BigQuery, Google Cloud Storage, and Pub/Sub. Reference: https://cloud.google.com/stackdriver/

Community vote distribution

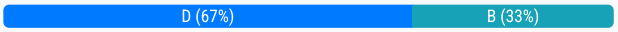

JoeShmoe Highly Voted 6 years, 9 months ago
D is correct and best practice for long term log storage
upvoted 152 times

JoeShmoe Highly Voted 6 years, 9 months ago
D is correct and best practice for long term log storage+1. For archival purposes, Customer should use Cloud Storage. BigQuery is a datawarehouse, and could eventually import data from Cloud Storage if necessary.+1 Due to long term storage, cloud storage is better answer than BigQuerysame for me. The best approach for the long time log is to export it from monitoring to Cloud Storage Archival type
upvoted 152 times

AndreUanKenobi 5 years, 5 months ago
+1. For archival purposes, Customer should use Cloud Storage. BigQuery is a datawarehouse, and could eventually import data from Cloud Storage if necessary.
upvoted 22 times

AndreUanKenobi 5 years, 5 months ago
+1. For archival purposes, Customer should use Cloud Storage. BigQuery is a datawarehouse, and could eventually import data from Cloud Storage if necessary.+1 Due to long term storage, cloud storage is better answer than BigQuery
upvoted 22 times

anjuagrawal 4 years, 7 months ago
+1 Due to long term storage, cloud storage is better answer than BigQuery
upvoted 14 times

AndreaMa 2 years, 1 month ago
same for me. The best approach for the long time log is to export it from monitoring to Cloud Storage Archival type
upvoted 1 times

MeasService 6 years, 7 months ago
A and C can be quickly ruled out because none of them is solution for the requirements "retained for 5 years"

Between B and D, the different is where to store, BigQuery or Cloud Storage. Since the main concern is extended storing period, D (Correct Answer) is better choice, and the "retained for 5 years for future analysis" further qualifies it, for example, using Coldline storage class.

With regards of BigQuery, while it is also a low-cost storage, but the main purpose is for analysis. Also, logs stored in Cloud Storage is easy to transport to BigQuery or do query directly against the files saved in Cloud Storage if and whenever needed.
upvoted 75 times

MeasService 6 years, 7 months ago
A and C can be quickly ruled out because none of them is solution for the requirements "retained for 5 years" Between B and D, the different is where to store, BigQuery or Cloud Storage. Since the main concern is extended storing period, D (Correct Answer) is better choice, and the "retained for 5 years for future analysis" further qualifies it, for example, using Coldline storage class. With regards of BigQuery, while it is also a low-cost storage,

but the main purpose is for analysis. Also, logs stored in Cloud Storage is easy to transport to BigQuery or do query directly against the files saved in Cloud Storage if and whenever needed.point : organization requires that metrics from all applications be retained for 5 yearsI mean answer : DIf you have 2 viable solutions (B&D), then always chose the one that is cost optimised - I chose DBigquery long term storage cost: $0.020 per GB Cloud Storage archive cost: $0,0012 per GB Only if metrics need less than 10 GB (free service part on Bigquery) then the correct solution will be B... But all metrics for all applications during more than 5 years... I think never will be the case :Dsecond that! like the way u explained..The question is about metrics, not logs. I'd go for B. See https://cloud.google.com/solutions/stackdriver-monitoring-metric-exportThis is a good example. thanks. But, we can easily change that implementation to dump the metrics to buckets to save lots of money. And, when talking about legal purpose, 1 hour interval may not be enough. You may have to keep more frequent metrics. So, only cold line or archive work for that purpose.

upvoted 75 times

Shyeom 6 years, 7 months ago

point : organization requires that metrics from all applications be retained for 5 years

upvoted 2 times

Shyeom 6 years, 7 months ago

point : organization requires that metrics from all applications be retained for 5 yearsI mean answer : D

upvoted 2 times

Shyeom 6 years, 7 months ago

I mean answer : D

upvoted 4 times

Cloudy_Apple_Juice 5 years, 10 months ago

If you have 2 viable solutions (B&D), then always chose the one that is cost optimised - I chose D

upvoted 5 times

Cloudy_Apple_Juice 5 years, 10 months ago

If you have 2 viable solutions (B&D), then always chose the one that is cost optimised - I chose DBigquery long term storage cost: $0.020 per GB Cloud Storage archive cost: $0,0012 per GB Only if metrics need less than 10 GB (free service part on Bigquery) then the correct solution will be B... But all metrics for all applications during more than 5 years... I think never will be the case :D

upvoted 5 times

jvale 4 years, 5 months ago

Bigquery long term storage cost: $0.020 per GB Cloud Storage archive cost: $0,0012 per GB Only if metrics need less than 10 GB (free service part on Bigquery) then the correct solution will be B... But all metrics for all applications during more than 5 years... I think never will be the case :D

upvoted 11 times

Vika 5 years, 5 months ago

second that! like the way u explained..

upvoted 1 times

trainor 5 years, 8 months ago

The question is about metrics, not logs. I'd go for B. See https://cloud.google.com/solutions/stackdriver-monitoring-metric-export

upvoted 18 times

trainor 5 years, 8 months ago

The question is about metrics, not logs. I'd go for B. See https://cloud.google.com/solutions/stackdriver-monitoring-metric-exportThis is a good example. thanks. But, we can easily change that implementation to dump the metrics to buckets to save lots of money. And, when talking about legal purpose, 1 hour interval may not be enough. You may have to keep more frequent metrics. So, only cold line or archive work for that purpose.

upvoted 18 times

bnlcnd 5 years, 6 months ago

This is a good example. thanks. But, we can easily change that implementation to dump the metrics to buckets to save lots of money. And, when talking about legal purpose, 1 hour interval may not be enough. You may have to keep more frequent metrics. So, only cold line or archive work for that purpose.

upvoted 5 times

bruh_1 Most Recent 7 months, 2 weeks ago

it says "future analysis" so BigQuery

upvoted 2 times

bigdapper 1 year ago

B. BigQuery is correct because the question specifically mentions future *analysis*.

upvoted 2 times

gaufchamp 1 year, 4 months ago

Stackdriver Monitoring (now known as Google Cloud Monitoring) allows you to collect metrics from all your applications and services. By exporting the metrics to BigQuery, you ensure long-term storage and analysis capability. BigQuery allows you to store vast amounts of data (such as metrics)

and retain it for extended periods, like 5 years, while providing powerful querying capabilities for future analysis, including in legal proceedings. BigQuery is ideal because it allows you to query the data at any time, and it is designed for long-term storage and analysis of large datasets.

upvoted 2 times

halifax 1 year, 5 months ago

the question specifically mentions "future analysis in possible legal proceedings." This makes BigQuery the most appropriate choice, even if it's more expensive. option D is cheaper (cloud storage), but data stored at cloud storage (object storage) can not be analysed.

upvoted 2 times

hpf97 1 year, 6 months ago

D : Stackdriver for logging, then Storage with data policy for moving data to less expensive areas. And data is restored in BigQuery only if legal investigation required.

upvoted 1 times

alihabib 1 year, 8 months ago

Since the data is not actively queried, and cited as "possible" use qualifies it to be archived. Which means Cloud Storage

upvoted 1 times

drinkwater 1 year, 8 months ago

B is the right answer why is not D, because while Google Cloud Storage can handle long-term storage, it is less efficient than BigQuery for analysis. Retrieving and querying metrics from Cloud Storage would require additional tools or steps, making it less suitable for the described use case.

upvoted 1 times

alpay 1 year, 9 months ago

B in my opinion. Take a look: https://cloud.google.com/architecture/monitoring-metric-export#store_metrics BQ is also long-term storage with option to reduce cost of older data. https://cloud.google.com/bigquery/docs/best-practices-storage

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Long-term Retention: BigQuery is a data warehouse designed for long-term storage and analysis of large datasets. It's the ideal place to store metrics for 5 years to meet your organization's legal requirements. 2. Cost-Effective: BigQuery's storage pricing is very competitive, especially for long-term data retention. 3. Analysis and Reporting: BigQuery provides powerful tools for analyzing and querying data, making it easy to extract insights and generate reports from the stored metrics. 4. Integration: Stackdriver Monitoring (now Cloud Monitoring) can be easily configured to export metrics to BigQuery. D is not correct as while Cloud Storage can store data for long periods, it's not optimized for querying and analyzing data like BigQuery.

upvoted 2 times

VedaSW 1 year, 11 months ago

I go for B, as the question is about 5 years worth of data "for future analysis in possible legal proceedings", and the "future" can be next day, based on when the legal proceeding happen. It is not about long term log storage. Even the argument of "future" means 100 years later, the Cold Storage Archival still does not fulfill the "analysis" portion of the requirements. You will need to move the data from Cold Storage to BigQuery for the analysis. So the ideal answer should be combination of D and B, but we do not have such option, hence the answer can meet all requirements is B.

upvoted 1 times

Hungdv 2 years ago

D is answer. Monitoring has only 24 months retention.

upvoted 1 times

joecloud12 2 years ago

Selected Answer: D

for storing the logs you need cloud storage.

upvoted 1 times

joecloud12 2 years ago

for storing the logs you need cloud storage.The question didn't mention logs.

upvoted 1 times

ryaryarya 1 year, 7 months ago

The question didn't mention logs.

upvoted 1 times

monus 2 years ago

Selected Answer: B

B should be correct. How can Cloud Storage analyze the data?

upvoted 2 times

monus 2 years ago

B should be correct. How can Cloud Storage analyze the data?The statement is not about the actual analysis of the data, but 'where' to store the data for future analysis. Who know when that will be??? So GCS is the best answer. When need be, it can be import into BQAlso it said retained for 5 years for future analysis... you don't store in BQ

upvoted 2 times

desertlotus1211 2 years ago

The statement is not about the actual analysis of the data, but 'where' to store the data for future analysis. Who know when that will be??? So GCS is the best answer. When need be, it can be import into BQAlso it said retained for 5 years for future analysis... you don't store in BQ

upvoted 1 times

desertlotus1211 2 years ago

The statement is not about the actual analysis of the data, but 'where' to store the data for future analysis. Who know when that will be??? So GCS is the best answer. When need be, it can be import into BQ

upvoted 1 times

desertlotus1211 2 years ago

Also it said retained for 5 years for future analysis... you don't store in BQ

upvoted 1 times

a2le 2 years, 2 months ago

I mean, "for possible future legal proceedings", I think that immutable storage that grants data integrity is the best option here, what's more, it's also the cheapest one...

upvoted 1 times

james2033 2 years, 2 months ago

Google Cloud Storage for 5 years storing legally.

upvoted 1 times

Question #27 *Topic 1*

Your company has decided to build a backup replica of their on-premises user authentication PostgreSQL database on Google Cloud Platform. The database is 4
TB, and large updates are frequent. Replication requires private address space communication.
Which networking approach should you use?

A. Google Cloud Dedicated Interconnect

B. Google Cloud VPN connected to the data center network

C. A NAT and TLS translation gateway installed on-premises

D. A Google Compute Engine instance with a VPN server installed connected to the data center network

Suggested Answer: *A*
Google Cloud Dedicated Interconnect provides direct physical connections and RFC 1918 communication between your on-premises network and Google's network. Dedicated Interconnect enables you to transfer large amounts of data between networks, which can be more cost effective than purchasing additional bandwidth over the public Internet or using VPN tunnels. Benefits: ✑ Traffic between your on-premises network and your VPC network doesn't traverse the public Internet. Traffic traverses a dedicated connection with fewer hops, meaning there are less points of failure where traffic might get dropped or disrupted. ✑ Your VPC network's internal (RFC 1918) IP addresses are directly accessible from your on-premises network. You don't need to use a NAT device or VPN tunnel to reach internal IP addresses. Currently, you can only reach internal IP addresses over a dedicated connection. To reach Google external IP addresses, you must use a separate connection. ✑ You can scale your connection to Google based on your needs. Connection capacity is delivered over one or more 10 Gbps Ethernet connections, with a maximum of eight connections (80 Gbps total per interconnect). ✑ The cost of egress traffic from your VPC network to your on-premises network is reduced. A dedicated connection is generally the least expensive method if you have a high-volume of traffic to and from Google's network. Reference: https://cloud.google.com/interconnect/docs/details/dedicated

Community vote distribution

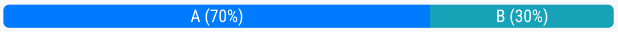

AWS56 Highly Voted 6 years, 7 months ago
A is the one
upvoted 29 times

AWS56 Highly Voted 6 years, 7 months ago
A is the oneA is okA, direct connect is private. VPN not enough for 4 TB with huge frequent changes.
upvoted 29 times

tartar 6 years ago
A is ok
upvoted 8 times

nitinz 5 years, 5 months ago
A, direct connect is private. VPN not enough for 4 TB with huge frequent changes.
upvoted 3 times

amxexam 4 years, 11 months ago
Let's go with option elimination A. Google Cloud Dedicated Interconnect >> Secured, fast connection, hence the choice. This will allow private connection from GCP to the data centre with a fast connection. Cost is not mentioned in the requirement to eliminate this option. B. Google Cloud VPN connected to the data centre network >> We have to think about data flowing on the internet and the requirement talks about private connect. Also not sure how well you connect VPN with Data Center until you use the hybrid option. https://cloud.google.com/network-connectivity/docs/vpn/concepts/overview hence eliminate C. A NAT and TLS translation gateway installed on-premises >>This is a VM option to reach outside won't for this requirement hence eliminate D. A Google Compute Engine instance with a VPN server installed connected to the data centre network >>This is a slow option hence eliminate Hence A
upvoted 19 times

gaufchamp Most Recent 1 year, 4 months ago
Google Cloud Dedicated Interconnect provides the necessary performance, reliability, and private connectivity required for your database replication scenario. Therefore, Option A is the best choice.
upvoted 1 times

alihabib 1 year, 6 months ago

A - Private Connection to VPC network, and transfer speed of 10 GBPS makes it obvious choice to maintain fast read replicas

upvoted 2 times

ramjisriram 1 year, 7 months ago

Key clue is "Replication requires private address space communication" . Only Google Cloud Dedicated Interconnect has private address. VPN is public, encrypted and slower !!! Another clue is "Large updates are frequent" means you need faster Dedicated interconnect, as VPN will be slower.

upvoted 3 times

Ekramy_Elnaggar 1 year, 9 months ago

1. High Bandwidth and Reliability: Dedicated Interconnect provides a direct physical connection between your on-premises network and Google Cloud, offering high bandwidth and low latency. This is essential for replicating a 4TB database with frequent large updates. 2. Private Address Space: Dedicated Interconnect allows you to extend your private IP address space to Google Cloud, ensuring secure and private communication for database replication. 3. Security: Dedicated Interconnect offers a more secure connection compared to VPN, as traffic doesn't traverse the public internet. Note: Question didn't mention anything about costs, so guys please stop overthinking and focus on the question key words.

upvoted 6 times

Robert0 2 years, 2 months ago

Option B Interconnect is incredibly expensive and the usecase do not justify it. Properly configured, a VPN provides similar features. If the question included an inusally high SLA, I will go with Interconnect. If not, VPN is a great option.A simple VPN may not provide enough bandwidth for replication if the DB is busy. We know that the auth DB is 4TB, I'd say this must be a quite big company, possibly they can offer an interconnect? But it surely is expensive

upvoted 5 times

Robert0 2 years, 2 months ago

Selected Answer: B

Option B Interconnect is incredibly expensive and the usecase do not justify it. Properly configured, a VPN provides similar features. If the question included an inusally high SLA, I will go with Interconnect. If not, VPN is a great option.

upvoted 5 times

tocsa 2 years, 2 months ago

A simple VPN may not provide enough bandwidth for replication if the DB is busy. We know that the auth DB is 4TB, I'd say this must be a quite big company, possibly they can offer an interconnect? But it surely is expensive

upvoted 1 times

sidiosidi 2 years, 3 months ago

Selected Answer: B

I'll go for VPN. First, the database is only for authentication and updates will be on this part, small portion of data needs to be replicated between on-premise and cloud. so no need for high bandwidth. the first migration will needs bandwidth but not toom much (5T) can be migrated using VPN. VPN permits to use private networking and it's secure. VPN not expensive as direct connect. As architect you should also evaluate the cost over the requirement, at the end you need convice business with solution. Paying 5K will kick you out the project for such small requirement.

upvoted 3 times

sidiosidi 2 years, 3 months ago

I'll go for VPN. First, the database is only for authentication and updates will be on this part, small portion of data needs to be replicated between on-premise and cloud. so no need for high bandwidth. the first migration will needs bandwidth but not toom much (5T) can be migrated using VPN. VPN permits to use private networking and it's secure. VPN not expensive as direct connect. As architect you should also evaluate the cost over the requirement, at the end you need convice business with solution. Paying 5K will kick you out the project for such small requirement.I do not agree, it is said that the 4TB database is for user authentication database; so it cannot be part of it.I think this fella hit the righ nail. Interconnect is incredibly expensive and the usecase do not justify it. Properly configured, a VPN provides similar features. If the question included an inusally high SLA, I will go with Interconnect. If not, VPN is a great option.

upvoted 3 times

hpf97 1 year, 6 months ago

I do not agree, it is said that the 4TB database is for user authentication database; so it cannot be part of it.

upvoted 1 times

Robert0 2 years, 2 months ago

I think this fella hit the righ nail. Interconnect is incredibly expensive and the usecase do not justify it. Properly configured, a VPN provides similar features. If the question included an inusally high SLA, I will go with Interconnect. If not, VPN is a great option.

upvoted 1 times

Jen3 2 years, 5 months ago

If you tried to sell me on Interconnect when all I needed was a VPN (meets bandwidth req, private address space, encryption of traffic possible), I would reach out to AWS for a quote...

upvoted 1 times

lisabisa 2 years, 10 months ago

GoogleVPN throughput is 3Gbps. It supports private IP connection and cheaper than DIrect Connection. Direct connect supports 8 * 10Gbps or 2*100Gbps. But too expensive for this

upvoted 3 times

eka_nostra 3 years ago

Connect to private space with high-speed bandwidth will go to A.

upvoted 3 times

mrhege 3 years, 3 months ago

B: Dedicated Interconnect would be a major overkill here and a quite expensive one as well. Requirements mention private _address space_, not private connection. Data over VPN is just as secure. Also there is no mention that a Google PoP would be available.
https://cloud.google.com/network-connectivity/docs/how-to/choose-product

upvoted 1 times

mohideenks 3 years, 8 months ago

A is the correct answer

upvoted 1 times

Mahmoud_E 3 years, 10 months ago

A is great but expensive for just a database DR but what can we do about that

upvoted 1 times

zr79 3 years, 10 months ago

VPN is not private, it is public but encrypted. Also, VPN is not suitable for large updates that happen frequently

upvoted 1 times

AzureDP900 3 years, 10 months ago

without any second thought A is right

upvoted 1 times

minmin2020 3 years, 10 months ago

A. Google Cloud Dedicated Interconnect - large updates and better security, however may not be the most cost effective choice

upvoted 1 times

Question #28 *Topic 1*

Auditors visit your teams every 12 months and ask to review all the Google Cloud Identity and Access Management (Cloud IAM) policy changes in the previous 12 months. You want to streamline and expedite the analysis and audit process.
What should you do?

A. Create custom Google Stackdriver alerts and send them to the auditor

B. Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor

C. Use cloud functions to transfer log entries to Google Cloud SQL and use ACLs and views to limit an auditor's view

D. Enable Google Cloud Storage (GCS) log export to audit logs into a GCS bucket and delegate access to the bucket

Suggested Answer: *D*
Official Answer: D

Community vote distribution

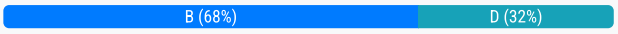

ghitesh Highly Voted 6 years, 7 months ago
B. https://cloud.google.com/iam/docs/roles-audit-logging#scenario_external_auditors
upvoted 100 times

ghitesh Highly Voted 6 years, 7 months ago
B. https://cloud.google.com/iam/docs/roles-audit-logging#scenario_external_auditorsb) seems correctI agree. Answer BThe article references either gcs or bq. I think this q is referring to gcsB makes more sense after reading it. thx
upvoted 100 times

rockstar9622 6 years, 7 months ago
b) seems correct
upvoted 3 times

anton_royce 6 years, 4 months ago
I agree. Answer B
upvoted 5 times

MikeB19 4 years, 11 months ago
The article references either gcs or bq. I think this q is referring to gcs
upvoted 1 times

TheCloudBoy77 4 years, 9 months ago
B makes more sense after reading it. thx
upvoted 4 times

jcmoranp 6 years, 10 months ago
Think B is better. Export to Bigquery and restrict access to queries with ACLs to auditorsD is ok.Sorry, changed my view. B is the recommended practicedon't change your view, D was right :)B is correctD, rest all options are no good.Please check the keywords in question -- "streamline and expedite" -- Bigquery is suitable not storage bucket. so it should be (B)I thought same as well. I would go with BI think D is better. B implies too much data manipulation to make it suitable for an audit.
upvoted 38 times

jcmoranp 6 years, 10 months ago
Think B is better. Export to Bigquery and restrict access to queries with ACLs to auditors
upvoted 38 times

tartar 6 years ago
D is ok.
upvoted 7 times

tartar 6 years ago
D is ok.Sorry, changed my view. B is the recommended practicedon't change your view, D was right :)B is correct
upvoted 7 times

tartar 6 years ago
Sorry, changed my view. B is the recommended practice

upvoted 14 times

tartar 6 years ago

Sorry, changed my view. B is the recommended practicedon't change your view, D was right :)B is correct

upvoted 14 times

alii 5 years, 7 months ago

don't change your view, D was right :)B is correct

upvoted 4 times

alii 5 years, 7 months ago

don't change your view, D was right :)

upvoted 4 times

RKS_2021 5 years, 1 month ago

B is correct

upvoted 1 times

nitinz 5 years, 5 months ago

D, rest all options are no good.

upvoted 3 times

nitinz 5 years, 5 months ago

D, rest all options are no good.Please check the keywords in question -- "streamline and expedite" -- Bigquery is suitable not storage bucket. so it should be (B)

upvoted 3 times

AmitAr 4 years, 3 months ago

Please check the keywords in question -- "streamline and expedite" -- Bigquery is suitable not storage bucket. so it should be (B)

upvoted 3 times

passnow 6 years, 8 months ago

I thought same as well. I would go with B

upvoted 5 times

trainor 5 years, 8 months ago

I think D is better. B implies too much data manipulation to make it suitable for an audit.

upvoted 4 times

ahmed75.zahran00 Most Recent 6 months, 2 weeks ago

expedite the analysis and audit process means B Minimize cost D

upvoted 1 times

hpf97 1 year, 6 months ago

A : not good as sending security data outside, even for audit, is not compliant C : cloud SQL is not designed for audits D : Would be great as the linly one to mention audit logs but only the log from GCS, and not IAM. B : remaining answer but it would be prefered audit log

upvoted 1 times

ramjisriram 1 year, 7 months ago

Key clue from the question is "You want to streamline and expedite the analysis". How can you expect streamline and analysis capabilities from Cloud Storage? Rather BigQuery has the analytics and streamline capabilities inbuilt in the tool. User access can be controlled in both, so the correct answer with key clues is to Enable Logging export to Google BigQuery (this will store all the logs streamlined) and use ACLs (these are filters, so according to the audior) and views (these are read only, there is no obfuscation or tokenization or sensistive data to be handled from the logs !!!!) to scope the data shared with the auditor.

upvoted 4 times

Ishu_awsguy 1 year, 7 months ago

The power word here is , Streamline analysis and audit , hence BQ is the answer. If it would have been only audit , then D would have been beteer

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Comprehensive Audit Trail: Cloud Logging automatically captures audit logs for all Cloud IAM activity. Exporting these logs to BigQuery provides a centralized and comprehensive audit trail for analysis. 2. Powerful Analysis: BigQuery's analytical capabilities allow auditors to efficiently query and analyze IAM policy changes over the 12-month period. They can filter, aggregate, and generate reports to identify any anomalies or security concerns. 3. Granular Access Control: BigQuery's Access Control Lists (ACLs) and views enable you to precisely control which data the auditors can access. This ensures that they only see the information relevant to their audit without exposing sensitive data. Note: While exporting logs to Cloud Storage is possible, it's less efficient for analysis compared to BigQuery.

upvoted 4 times

nareshthumma 1 year, 9 months ago

Answer B

upvoted 1 times

maxdanny 1 year, 11 months ago

Option B is the best approach. Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor. This method provides robust querying capabilities, ensures that historical IAM policy changes can be analyzed effectively, and allows you to control access securely.

upvoted 1 times

joecloud12 2 years ago

b is correct because it is easier to implement compared to D

upvoted 1 times

H_S 2 years, 1 month ago

READ THIS, ACL is not available in BIG QUERY , thereforeD. Enable Google Cloud Storage (GCS) log export to audit logs into a GCS bucket and delegate access to the bucket

upvoted 4 times

Jen3 2 years, 5 months ago

ACLs would provide year-round access to the data which is more privileges than necessary. Logs will need to be retained for a full year because hypothetically, January logs could be looked at in December. Cloud Storage offers signed URLs, and less expensive storage options.

upvoted 1 times

lisabisa 2 years, 5 months ago

Both B and D are ok. Using cloud storage requires additional setup for auditors, pulling data to BQ. Using BQ would satisfy "streamline and expedite the analysis and audit process"

upvoted 1 times

Teckexam 2 years, 7 months ago

Based on google documentation B is the correct answer. https://cloud.google.com/iam/docs/job-functions/auditing#scenario_external_auditors Dashboard is available in BigQuery to review historic logs and in case anamoly is found elevated access is provided. Access is revoked after audit activities are done.

upvoted 4 times

kip21 2 years, 7 months ago

D - Correct B - his option requires additional work to set up the ACLs and views to limit an auditor's view of the data. This could be time-consuming and complex to implement. Furthermore, BigQuery may not be the ideal tool for auditors who are only interested in reviewing Cloud IAM policy changes.

upvoted 2 times

CloudDom 2 years, 9 months ago

That‘s the only logical one also Bard is confirming this one

upvoted 1 times

thewalker 2 years, 9 months ago

D I will not go with B, as the requirement is once for 12 months. Push the data in Coldline for 12 months and retrieve it during audit is enough. Save costs.

upvoted 3 times

thewalker 2 years, 9 months ago

D I will not go with B, as the requirement is once for 12 months. Push the data in Coldline for 12 months and retrieve it during audit is enough. Save costs.Coldline / ArchiveStreamline and expedite analysis is the goal. Costs are never brought up.

upvoted 3 times

thewalker 2 years, 9 months ago

Coldline / Archive

upvoted 2 times

hogtrough 2 years, 7 months ago

Streamline and expedite analysis is the goal. Costs are never brought up.

upvoted 2 times

Question #29 *Topic 1*

You are designing a large distributed application with 30 microservices. Each of your distributed microservices needs to connect to a database back-end. You want to store the credentials securely.
Where should you store the credentials?

A. In the source code

B. In an environment variable

C. In a secret management system

D. In a config file that has restricted access through ACLs

Suggested Answer: *C*
Reference: https://cloud.google.com/kms/docs/secret-management

Community vote distribution
C (100%)

Eroc Highly Voted 6 years, 9 months ago
Google Secret Management was designed explicitly for this purpose.
upvoted 34 times

Eroc Highly Voted 6 years, 9 months ago
Google Secret Management was designed explicitly for this purpose.C is okC, microservices = GKE = Kubernetes = secrets.
upvoted 34 times

tartar 6 years ago
C is ok
upvoted 7 times

nitinz 5 years, 5 months ago
C, microservices = GKE = Kubernetes = secrets.
upvoted 6 times

shandy 6 years, 8 months ago
C is the answer, since key management systems generate, use, rotate, encrypt, and destroy cryptographic keys and manage permissions to those keys. A is incorrect because storing credentials in source code and source control is discoverable, in plain text, by anyone with access to the source code. This also introduces the requirement to update code and do a deployment each time the credentials are rotated. B is not correct because consistently populating environment variables would require the credentials to be available, in plain text, when the session is started. D is incorrect because instead of managing access to the config file and updating manually as keys are rotated, it would be better to leverage a key management system. Additionally, there is increased risk if the config file contains the credentials in plain text.
upvoted 14 times

Ekramy_Elnaggar Most Recent 1 year, 9 months ago
1. Centralized and Secure Storage: Secret management systems like HashiCorp Vault, AWS Secrets Manager, or Google Cloud Secret Manager provide a centralized and secure location to store sensitive credentials. This ensures that database credentials are not scattered across multiple microservices or configuration files. 2. Access Control: Secret management systems offer fine-grained access control, allowing you to restrict access to secrets based on roles and permissions. This ensures that only authorized microservices and users can access the database credentials. 3. Rotation and Auditing: These systems often provide features for automatic secret rotation and auditing, which helps improve security and compliance. 4. Integration: Secret management systems can integrate with your deployment pipelines and orchestration tools, making it easier to manage secrets throughout the application lifecycle.
upvoted 3 times

ashishdwi007 2 years, 6 months ago
Other options are not best practices.
upvoted 1 times

Teckexam 2 years, 7 months ago
Need to use key management system for this usecase since other options are not secure.
upvoted 1 times

hiromi 3 years, 4 months ago

C is correct

upvoted 1 times

omermahgoub 3 years, 8 months ago

C. In a secret management system It is important to store the credentials for your database back-end securely in order to protect them from unauthorized access. One way to do this is by using a secret management system, such as Google Cloud's Secret Manager. Secret Manager is a secure and convenient storage system for API keys, passwords, and other sensitive data that is designed to protect against unauthorized access. By storing the credentials in Secret Manager, you can ensure that they are kept secure and can be easily accessed by your microservices as needed. Storing the credentials in the source code, an environment variable, or a config file that has restricted access through ACLs may not provide the same level of security as a dedicated secret management system. It is important to ensure that your credentials are stored in a secure and controlled manner to protect against unauthorized access.

upvoted 4 times

AniketD 3 years, 9 months ago

C is correct; If credential then always use secret manager.

upvoted 1 times

megumin 3 years, 9 months ago

C is ok

upvoted 1 times

zr79 3 years, 10 months ago

secret manager is the answer

upvoted 1 times

AzureDP900 3 years, 10 months ago

C is right

upvoted 1 times

Kubernetes 3 years, 11 months ago

answer is C

upvoted 1 times

rogerlovato 4 years, 7 months ago

C is correct

upvoted 1 times

haroldbenites 4 years, 8 months ago

Go for C

upvoted 1 times

vincy2202 4 years, 8 months ago

C is the right answer

upvoted 1 times

unnikrisb 4 years, 10 months ago

Google Practice exam question with option C : In a key management system C is correct because key management systems generate, use, rotate, encrypt, and destroy cryptographic keys and manage permissions to those keys. https://cloud.google.com/kms/ For this question, refer to the Mountkirk Games case study.

upvoted 1 times

unnikrisb 4 years, 10 months ago

Google Practice exam question with option C : In a key management system Here also C is correct because key management systems generate, use, rotate, encrypt, and destroy cryptographic keys and manage permissions to those keys. https://cloud.google.com/kms/ For this question, refer to the Mountkirk Games case study.

upvoted 1 times

Question #30 *Topic 1*

A lead engineer wrote a custom tool that deploys virtual machines in the legacy data center. He wants to migrate the custom tool to the new cloud environment.
You want to advocate for the adoption of Google Cloud Deployment Manager.
What are two business risks of migrating to Cloud Deployment Manager? (Choose two.)

A. Cloud Deployment Manager uses Python

B. Cloud Deployment Manager APIs could be deprecated in the future

C. Cloud Deployment Manager is unfamiliar to the company's engineers

D. Cloud Deployment Manager requires a Google APIs service account to run

E. Cloud Deployment Manager can be used to permanently delete cloud resources

F. Cloud Deployment Manager only supports automation of Google Cloud resources

Suggested Answer: *B, F*
Official Answer: BF

Community vote distribution


victory108 Highly Voted 5 years, 1 month ago
E. Cloud Deployment Manager can be used to permanently delete cloud resources F. Cloud Deployment Manager only supports automation of Google Cloud resourcesYup, E + F. In GCP documentation it states as a warning note that deletion made through Deployment Manager scripts cannot be undone, if devs are not well trained a human errors can impact BusinessF is not a risk, it is a limitation of solution. Risk is something that is not known for sure and is manageable (risk can be mitigated, avoided). You cannot manage the limitation of solution. You can use it with this limitation or not and you know it in advance.Advocating to adopt Google Cloud Deployment Manager will become a risk if the lead engineer or other business need ask for use other cloud platform.
upvoted 80 times

victory108 Highly Voted 5 years, 1 month ago
E. Cloud Deployment Manager can be used to permanently delete cloud resources F. Cloud Deployment Manager only supports automation of Google Cloud resources
upvoted 80 times

poseidon24 5 years ago
Yup, E + F. In GCP documentation it states as a warning note that deletion made through Deployment Manager scripts cannot be undone, if devs are not well trained a human errors can impact Business
upvoted 14 times

Gregwaw 2 years, 10 months ago
F is not a risk, it is a limitation of solution. Risk is something that is not known for sure and is manageable (risk can be mitigated, avoided). You cannot manage the limitation of solution. You can use it with this limitation or not and you know it in advance.Advocating to adopt Google Cloud Deployment Manager will become a risk if the lead engineer or other business need ask for use other cloud platform.
upvoted 5 times

Gregwaw 2 years, 10 months ago
F is not a risk, it is a limitation of solution. Risk is something that is not known for sure and is manageable (risk can be mitigated, avoided). You cannot manage the limitation of solution. You can use it with this limitation or not and you know it in advance.
upvoted 5 times

Terryhsieh 2 years, 7 months ago
Advocating to adopt Google Cloud Deployment Manager will become a risk if the lead engineer or other business need ask for use other cloud platform.
upvoted 3 times

AK2020 5 years, 2 months ago
C and F- make sense to mel think this is right. the key of the question is "business risks".yes, C and F right
upvoted 43 times

AK2020 5 years, 2 months ago

C and F- make sense to me

upvoted 43 times

ssepiro 4 years, 9 months ago

I think this is right. the key of the question is "business risks".

upvoted 3 times

ssepiro 4 years, 9 months ago

I think this is right. the key of the question is "business risks".yes, C and F right

upvoted 3 times

AzureDP900 3 years, 10 months ago

yes, C and F right

upvoted 3 times

marmar11111 Most Recent 10 months ago

E. Cloud Deployment Manager can be used to permanently delete cloud resources F. Cloud Deployment Manager only supports automation of Google Cloud resources

upvoted 1 times

Beetlesnail 1 year, 2 months ago

Going against the grain here but I strongly believe it's B and C. B. addresses long-term strategic and financial risks from vendor dependency C. addresses immediate operational and cost risks from knowledge gaps Why I think the others are wrong: A. Uses Python: Python is widely adopted in the industry, making this more of a minor technical consideration than a business risk. D. Requires service accounts: This is standard security practice and represents good architectural design rather than a business risk. E. Can delete resources: While this has operational implications, it's true of most infrastructure tools and can be mitigated through proper access controls and procedures. F. Google Cloud only: While this could represent vendor lock-in, since they're already migrating to Google Cloud environment, this limitation aligns with their current strategy.

upvoted 3 times

keyik777 1 year, 3 months ago

B: It will be deprecated on December 31, 2025 F: the tool is written for on prem data center

upvoted 2 times

ktan13 1 year, 5 months ago

These are business risk

upvoted 1 times

izekc 1 year, 5 months ago

CE looks more possible and relavent to risk

upvoted 2 times

hpf97 1 year, 6 months ago

So the aim solution should turn on "Business risk" A : no technical requirement B : as any COTS C : yes, so setting up the system can take time, and time is money, so business risk D : technical requirement, as for other IaC tool E : if the system is definitevely deleted, money is lost and bad reputation F : that's a fact, a limitation Futhermore when there is risk actions could be taken to mitigate them C : train people, hire experience ones E : code reviews, CI/CD... So C&E is the good couple

upvoted 2 times

user263263 1 year, 6 months ago

Definition of risk: "the possibility of something bad happening" A. is a fact B. Cloud Deployment Manager APIs could be deprecated in the future - part of it was already deprecated, see https://cloud.google.com/deployment-manager/docs/deprecations/composite-types, Deprecation is always a business risk for cloud services. There is some probability and if it happens, it will cost time and money to resolve. C. if it's true, it's a fact D. that's false, see example deployment without service account https://cloud.google.com/deployment-manager/docs/manage-cloud-resources-deployment E. Cloud Deployment Manager can be used to permanently delete cloud resources - there's a probability (hint: "can") that this causes some bad things to happen F. maybe that was a fact when the question was written, using 3rd party APIs is in beta now https://cloud.google.com/deployment-manager/docs/configuration/type-providers/process-adding-api

upvoted 2 times

LeoSantos121212121212121 1 year, 6 months ago

These are the business risks

upvoted 2 times

kino_1994 1 year, 7 months ago

My answer is BF and aligns with similar arguments as those presented by @Choopaower B: This option can be deprecated, as the custom tool has been used for a long time. F: Cloud Deployment Manager only supports automation for Google Cloud resources, not custom resources as often happens in On-Premise environments. C is not a business risk. It is a technical risk. Their engineers may not know it yet, but they can learn.

upvoted 1 times

ramjisriram 1 year, 7 months ago

B and C are the business risks, rest of the options are technical limitations. B. Cloud Deployment Manager APIs could be deprecated in the future C. Cloud Deployment Manager is unfamiliar to the company's engineers The other options might present technical challenges or considerations but are not necessarily business risks. For example: A is more about the programming language used. D is about the service account requirement. E is a capability that, if misused, could cause issues but is not inherently a business risk. F limits scope but isn't a direct business risk.

upvoted 3 times

drinkwater 1 year, 8 months ago

the right answers are C & E C: because the engineers are dealing with legacy technologies E: Cloud Deployment Manager has the ability to delete resources

upvoted 2 times

Ekramy_Elnaggar 1 year, 9 months ago

I guess everyone agree on C, the debate is between B and F. Before I start, I need to stress on the word "Business" not "Technical" risks, we need to take this in our minds. B: Cloud Deployment Manager APIs could be deprecated in the future >> this is not a risk at all as the APIs cannot be changed without an announcement and also there is a backward compatibility to prevent such issues, this is something that is tackled a decade ago in all cloud providers. F: Cloud Deployment Manager only supports automation of Google Cloud >> Indeed this is a business risk, as the question didn't mention explicitly that the customer is using only GCP, they left it vague by saying "new cloud environment", so we have to assume that it is multi-cloud strategy including On-premise BTW (Hybrid-Cloud as well), so we need a tool that can handle deployments on all of those targets. Based on all of that, I recommend (C & F)

upvoted 7 times

beagle_Masato 1 year, 9 months ago

C and F right

upvoted 3 times

Shasha1 1 year, 10 months ago

BF are correct

upvoted 1 times

Leo212003 1 year, 10 months ago

E and F

upvoted 1 times

Question #31 *Topic 1*

A development manager is building a new application. He asks you to review his requirements and identify what cloud technologies he can use to meet them. The application must:
1. Be based on open-source technology for cloud portability
2. Dynamically scale compute capacity based on demand
3. Support continuous software delivery
4. Run multiple segregated copies of the same application stack
5. Deploy application bundles using dynamic templates
6. Route network traffic to specific services based on URL
Which combination of technologies will meet all of his requirements?

A. Google Kubernetes Engine, Jenkins, and Helm

B. Google Kubernetes Engine and Cloud Load Balancing

C. Google Kubernetes Engine and Cloud Deployment Manager

D. Google Kubernetes Engine, Jenkins, and Cloud Load Balancing

Suggested Answer: *D*

▪

▪

Jenkins is an open-source automation server that lets you flexibly orchestrate your build, test, and deployment pipelines. Kubernetes Engine is a hosted version of Kubernetes, a powerful cluster manager and orchestration system for containers. When you need to set up a continuous delivery (CD) pipeline, deploying Jenkins on Kubernetes Engine provides important benefits over a standard VM-based deployment Incorrect Answers: A: Helm is a tool for managing Kubernetes charts. Charts are packages of pre-configured Kubernetes resources. Use Helm to: Find and use popular software packaged as Kubernetes charts

➪ Share your own applications as Kubernetes charts ➪ Create reproducible builds of your Kubernetes applications ➪ Intelligently manage your Kubernetes manifest files ➪ Manage releases of Helm packages Reference: https://cloud.google.com/solutions/jenkins-on-kubernetes-engine

Community vote distribution



rsamant Highly Voted 5 years, 2 months ago
it should be A .. helm is needed for "Deploy application bundles using dynamic templates"

Load Balancing should be part of GKE Already
upvoted 70 times

rsamant Highly Voted 5 years, 2 months ago
it should be A .. helm is needed for "Deploy application bundles using dynamic templates" Load Balancing should be part of GKE AlreadyA should be fineLoad balancing is not a part of GKE untill it's created explicitlyKubernetes Engine offers integrated support for two types of Cloud Load Balancing (Ingress and External Network Load Balancing) , hence Option A Reference : https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancerNot for "based on URL", that is the difference.https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer#optional_serving_multiple_applications_on_a_load_balancer As per the above document and given example of "fanout-ingress.yaml" in above document and also in GKE sample repository below https://github.com/GoogleCloudPlatform/kubernetes-engine-samples/tree/master/load-balancing it's clear that GKE LB can handle "6. Route network traffic to specific services based on URL" So NO need for Cloud Load balancing. Helm satisfy "5. Deploy application bundles using dynamic templates" and no other option satisfies this point #5. So correct answer should be: A
upvoted 70 times

AzureDP900 3 years, 10 months ago
A should be fine
upvoted 3 times

Prakzz 2 years, 10 months ago

Load balancing is not a part of GKE untill it's created explicitly

upvoted 2 times

raf2121 4 years, 11 months ago

Kubernetes Engine offers integrated support for two types of Cloud Load Balancing (Ingress and External Network Load Balancing) , hence Option A
Reference : https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer

upvoted 4 times

poseidon24 5 years ago

Not for "based on URL", that is the difference.

upvoted 11 times

poseidon24 5 years ago

Not for "based on URL", that is the difference.https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer#optional_serving_multiple_applications_on_a_load_balancer As per the above document and given example of "fanout-ingress.yaml" in above document and also in GKE sample repository below https://github.com/GoogleCloudPlatform/kubernetes-engine-samples/tree/master/load-balancing it's clear that GKE LB can handle "6. Route network traffic to specific services based on URL" So NO need for Cloud Load balancing. Helm satisfy "5. Deploy application bundles using dynamic templates" and no other option satisfies this point #5. So correct answer should be: A

upvoted 11 times

ashish_t 4 years, 10 months ago

https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer#optional_serving_multiple_applications_on_a_load_balancer As per the above document and given example of "fanout-ingress.yaml" in above document and also in GKE sample repository below https://github.com/GoogleCloudPlatform/kubernetes-engine-samples/tree/master/load-balancing it's clear that GKE LB can handle "6. Route network traffic to specific services based on URL" So NO need for Cloud Load balancing. Helm satisfy "5. Deploy application bundles using dynamic templates" and no other option satisfies this point #5. So correct answer should be: A

upvoted 17 times

victory108 5 years, 1 month ago

D. Google Kubernetes Engine, Jenkins, and Cloud Load BalancingCloud Load Balancing does not fulfill "Be based on open-source technology for cloud portability"

upvoted 43 times

victory108 5 years, 1 month ago

D. Google Kubernetes Engine, Jenkins, and Cloud Load Balancing

upvoted 43 times

user263263 1 year, 6 months ago

Cloud Load Balancing does not fulfill "Be based on open-source technology for cloud portability"

upvoted 1 times

JoseCloudEng1994 Most Recent 11 months, 3 weeks ago

The answer should be A. Kubernetes and Helm are open-source but they are pointing out that a requirement is portability. Cload Load Balancing and Deployment Manager are GCP specific. So the only option remaining is A

upvoted 1 times

bigdapper 1 year ago

A is the only answer that contains all items based on open-source technology for cloud portability.

upvoted 1 times

izekc 1 year, 5 months ago

For networking part, D is required

upvoted 2 times

cloud_rider 1 year, 5 months ago

A is correct

upvoted 1 times

halifax 1 year, 5 months ago

Option A (Google Kubernetes Engine, Jenkins, and Helm): This option explicitly includes Helm, which is used for templated deployments. However, since GKE itself supports Helm, this option may be seen as redundant in the context of the requirements. Option D (Google Kubernetes Engine, Jenkins, and Cloud Load Balancing): This option includes Jenkins for continuous delivery and Cloud Load Balancing for routing traffic. While it doesn't mention Helm, the dynamic template deployment can still be managed through Jenkins scripts or other CI/CD tools.

upvoted 1 times

MikeMike7 1 year, 8 months ago

D: load balancing is needed for this requirement: 6. Route network traffic to specific services based on URL

upvoted 1 times

drinkwater 1 year, 8 months ago

A & D are both right Why D Might Still Be Preferred: While A is a valid choice, D (Google Kubernetes Engine, Jenkins, and Cloud Load Balancing) might be preferred for the following reasons: Cloud Load Balancing provides a more feature-rich, fully managed solution for routing traffic across multiple regions and services, including advanced load balancing, SSL termination, and support for more sophisticated network traffic management. It integrates well with GKE and offers additional scalability and flexibility that might be important as your system grows

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Open-source technology: Kubernetes Engine, Jenkins, and Helm. 2. Dynamically scale compute capacity: Kubernetes Engine provides autoscaling to adjust the number of nodes based on demand. 3. Support continuous software delivery: Jenkins enables CI/CD pipelines for automated building, testing, and deployment of applications. 4. Run multiple segregated copies: Kubernetes Engine allows deploying multiple instances of the application in isolated environments (namespaces) within the same cluster. 5. Deploy application bundles using dynamic templates: Helm uses charts (templates) to define, install, and upgrade Kubernetes applications. 6. Route network traffic based on URL: Kubernetes Engine's service objects and ingress controllers can route traffic to specific services based on URLs and other criteria.

upvoted 11 times

VedaSW 1 year, 11 months ago

Based on: "Be based on open-source technology for cloud portability", I go for A, because it is more portable, as compare to D.

upvoted 4 times

Hungdv 2 years ago

Choose A

upvoted 3 times

desertlotus1211 2 years ago

Yes, Kubernetes can route network traffic to specific services based on URL using Ingress and Ingress controllers Answer is A

upvoted 2 times

Chris_21 2 years, 1 month ago

Selected Answer: D

Load Balancer is required as per point 6. Jenkins satisfies point 3. D is correct

upvoted 1 times

Chris_21 2 years, 1 month ago

Load Balancer is required as per point 6. Jenkins satisfies point 3. D is correctHelm is needed for the dynamic templates requirement. The question is vague in whether the application is internally or externally facing which would clarify things a lot more for us. However, in its ambiguity option A has the techonologies neede to the requirments and thus deduce or infer that the application is internally facing and the ingress controllers will handle the routing of traffic. It's ambiguous on purpose which I hate in these exams. More of a test of how well we can read and interpret the questions vs our knowledge of material.

upvoted 1 times

Ric350 2 years ago

Helm is needed for the dynamic templates requirement. The question is vague in whether the application is internally or externally facing which would clarify things a lot more for us. However, in its ambiguity option A has the techonologies neede to the requirments and thus deduce or infer that the application is internally facing and the ingress controllers will handle the routing of traffic. It's ambiguous on purpose which I hate in these exams. More of a test of how well we can read and interpret the questions vs our knowledge of material.

upvoted 3 times

Rehamss 2 years, 5 months ago

D is correct

upvoted 1 times

kahinah 2 years, 5 months ago

Selected Answer: D

Option A (GKE, Jenkins, Helm) meets most requirements except for explicit URL-based routing, though Kubernetes Ingress (which can be managed through Helm charts) implicitly covers this.

Option D (GKE, Jenkins, Cloud Load Balancing) directly meets every requirement, including URL-based routing without needing to infer capabilities or integrate additional tools beyond the scope of what's listed. Jenkins supports continuous delivery, GKE supports dynamic scaling, segregated application stacks, and cloud portability. Cloud Load Balancing directly addresses the URL-based routing requirement.

upvoted 5 times

kahinah 2 years, 5 months ago

Option A (GKE, Jenkins, Helm) meets most requirements except for explicit URL-based routing, though Kubernetes Ingress (which can be managed through Helm charts) implicitly covers this. Option D (GKE, Jenkins, Cloud Load Balancing) directly meets every requirement, including URL-based

routing without needing to infer capabilities or integrate additional tools beyond the scope of what's listed. Jenkins supports continuous delivery, GKE supports dynamic scaling, segregated application stacks, and cloud portability. Cloud Load Balancing directly addresses the URL-based routing requirement.Except none of that meets the needs for deployment templates.My point exactly and my response to Chris_21. By process of elimination, you need Helm for the dynamic templates and you need Jenkins. Thus, you have to assume the application is internally facing and the ingress controllers will handle the traffic just fine.

upvoted 5 times

Sephethus 2 years, 2 months ago

Except none of that meets the needs for deployment templates.My point exactly and my response to Chris_21. By process of elimination, you need Helm for the dynamic templates and you need Jenkins. Thus, you have to assume the application is internally facing and the ingress controllers will handle the traffic just fine.

upvoted 1 times

Sephethus 2 years, 2 months ago

Except none of that meets the needs for deployment templates.

upvoted 1 times

Ric350 2 years ago

My point exactly and my response to Chris_21. By process of elimination, you need Helm for the dynamic templates and you need Jenkins. Thus, you have to assume the application is internally facing and the ingress controllers will handle the traffic just fine.

upvoted 1 times

VidhyaBupesh 2 years, 5 months ago

D is OK

upvoted 1 times

Question #32 *Topic 1*

You have created several pre-emptible Linux virtual machine instances using Google Compute Engine. You want to properly shut down your application before the virtual machines are preempted.
What should you do?

A. Create a shutdown script named k99.shutdown in the /etc/rc.6.d/ directory

B. Create a shutdown script registered as a xinetd service in Linux and configure a Stackdriver endpoint check to call the service

C. Create a shutdown script and use it as the value for a new metadata entry with the key shutdown-script in the Cloud Platform Console when you create the new virtual machine instance

D. Create a shutdown script, registered as a xinetd service in Linux, and use the gcloud compute instances add-metadata command to specify the service URL as the value for a new metadata entry with the key shutdown-script-url

Suggested Answer: *C*
A startup script, or a shutdown script, is specified through the metadata server, using startup script metadata keys. Reference: https://cloud.google.com/compute/docs/startupscript

Community vote distribution
C (65%) D (35%)

Eroc Highly Voted 6 years, 9 months ago
https://cloud.google.com/compute/docs/shutdownscript ... So CC, statup/shutdown script = metadataSince the instance is already created Option C gets eliminated. "gcloud compute instances addmetadata" command can be used to add or update the metadata of a virtual machine instance"
upvoted 43 times

Eroc Highly Voted 6 years, 9 months ago
https://cloud.google.com/compute/docs/shutdownscript ... So C
upvoted 43 times

nitinz 5 years, 5 months ago
C, statup/shutdown script = metadataSince the instance is already created Option C gets eliminated. "gcloud compute instances addmetadata" command can be used to add or update the metadata of a virtual machine instance"
upvoted 4 times

nitinz 5 years, 5 months ago
C, statup/shutdown script = metadata
upvoted 4 times

VishalB 5 years ago
Since the instance is already created Option C gets eliminated. "gcloud compute instances addmetadata" command can be used to add or update the metadata of a virtual machine instance"
upvoted 12 times

Gini 6 years, 3 months ago
I have doubts with the answer C because the question states that "You have created the instances" so C works too but the solution cannot apply to the already created instances. D seems correct to me... Reference: https://cloud.google.com/compute/docs/shutdownscript#apply_a_shutdown_script_to_running_instancesI also feel so because the virtual machines are already created.xinetd. Xinet makes the D answer be nonsenseYes. The correct answer should be D. To add a shutdown script to a running instance, follow the instructions in the Applying a startup script to running instances documentation but replace the metadata keys with one of the following keys: shutdown-script: Supply the shutdown script contents directly with this key. Using the gcloud command-line tool, you can provide the path to a shutdown script file, using the --metadata-from-file flag and the shutdown-script metadata key. shutdown-script-url: Supply a Cloud Storage URL to the shutdown script file with this key.changed my mind. preemptible vms can be stopped and started anytime. with that flexibility, C is ok.
upvoted 28 times

Gini 6 years, 3 months ago
I have doubts with the answer C because the question states that "You have created the instances" so C works too but the solution cannot apply to the already created instances. D seems correct to me...

Reference: https://cloud.google.com/compute/docs/shutdownscript#apply_a_shutdown_script_to_running_instances

upvoted 28 times

NG123 4 years, 2 months ago

I also feel so because the virtual machines are already created.

upvoted 2 times

dsnaghxhinwtsvvmip 3 years, 4 months ago

xinetd. Xinet makes the D answer be nonsense

upvoted 5 times

pepYash 5 years, 9 months ago

Yes. The correct answer should be D.

To add a shutdown script to a running instance, follow the instructions in the Applying a startup script to running instances documentation but replace the metadata keys with one of the following keys:

shutdown-script: Supply the shutdown script contents directly with this key. Using the gcloud command-line tool, you can provide the path to a shutdown script file, using the --metadata-from-file flag and the shutdown-script metadata key. shutdown-script-url: Supply a Cloud Storage URL to the shutdown script file with this key.

upvoted 4 times

pepYash 5 years, 9 months ago

Yes. The correct answer should be D. To add a shutdown script to a running instance, follow the instructions in the Applying a startup script to running instances documentation but replace the metadata keys with one of the following keys: shutdown-script: Supply the shutdown script contents directly with this key. Using the gcloud command-line tool, you can provide the path to a shutdown script file, using the --metadata-from-file flag and the shutdown-script metadata key. shutdown-script-url: Supply a Cloud Storage URL to the shutdown script file with this key.changed my mind. preemptible vms can be stopped and started anytime. with that flexibility, C is ok.

upvoted 4 times

pepYash 5 years, 9 months ago

changed my mind. preemptible vms can be stopped and started anytime. with that flexibility, C is ok.

upvoted 6 times

rajrocks171 Most Recent 1 year, 2 months ago

You'll typically use the shutdown-script or gce-shutdown-script metadata key for this purpose.

upvoted 1 times

halifax 1 year, 5 months ago

Option C is still a viable option even if the Linux VM instances are already created. You can add a shutdown script to existing instances using metadata. Option D is very complicated, The use of xinetd adds unnecessary complexity, and it may not be as effective as simply using the shutdown-script metadata key.

upvoted 1 times

RobertArnaud 1 year, 8 months ago

"when you create the new virtual machine instance" means it is too late, then C seems disqualified :(but the other choices are worse ?

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Preemptible instances and shutdown scripts: Google Cloud Platform's preemptible instances are cost-effective but can be terminated with short notice. To gracefully handle this, you need a shutdown script that runs before the instance is preempted. 2. Metadata and shutdown-script: Google Cloud allows you to add custom metadata to your instances. When you create a preemptible instance, you can include a shutdown-script metadata entry. The value of this entry should be your script, which will be executed automatically before the instance is preempted.

upvoted 1 times

Barry123456 1 year, 10 months ago

It's not C. C says "when you create the new virtual machine instance". But in the question, the instances have already been created. Thus, D. In D, "service URL" obviously means cloud storage.

upvoted 1 times

snehaso 2 years ago

Among Option C & D, option D uses shutdown-script-url but shutdown-script-url should be a Cloud storage url and not a linux service URL. That brings us to option C

upvoted 1 times

Hungdv 2 years ago

Choose C

upvoted 1 times

nicksb19 2 years, 1 month ago

C is correct since xinetd does not make sense.

upvoted 1 times

lisabisa 2 years, 5 months ago

Every virtual machine instance in GCP has access to a metadata server, which provides information about the instance and allows you to configure various settings, including startup and shutdown scripts. Startup and shutdown scripts are specified using special metadata keys in the metadata server. shutdown-script specifies the shutdown script that should be executed when the instance is being shut down.

upvoted 2 times

ashishdwi007 2 years, 6 months ago

All other options are related to play with Linux files or services. With Preemptible VMs ,these operations are overhead. Hence it makes sense to Automate such tasks.

upvoted 1 times

Mo7y 2 years, 7 months ago

The answer is either C or D I exclude D because shutdown-script-url only works with a shutdown script hosted on a cloud storage. Option D wants you to use shutdown-script-url for a locally hosted shutdown script, thus it's not the correct answer.

upvoted 3 times

kip21 2 years, 7 months ago

C https://cloud.google.com/compute/docs/shutdownscript

upvoted 1 times

Terryhsieh 2 years, 7 months ago

The answer should be C. reference to https://cloud.google.com/compute/docs/shutdownscript#apply_a_shutdown_script_to_running_instances
Regarding the answer D, it is not the option becasue no need to touch xinetd servie inside Linux.

upvoted 2 times

RLsh 2 years, 9 months ago

I believe the answer should be D since the VMs are already created

upvoted 1 times

Arun_m_123 2 years, 10 months ago

C is the right answer. See, there is one tip. In GCP, things like these are given to the customers as a solution - like give a shutdown script. GCP won't trouble the users to know all those geeky linux stuffs. So the answer is simply C

upvoted 1 times

Question #33 *Topic 1*

Your organization has a 3-tier web application deployed in the same network on Google Cloud Platform. Each tier (web, API, and database) scales independently of the others. Network traffic should flow through the web to the API tier and then on to the database tier. Traffic should not flow between the web and the database tier.
How should you configure the network?

A. Add each tier to a different subnetwork

B. Set up software based firewalls on individual VMs

C. Add tags to each tier and set up routes to allow the desired traffic flow

D. Add tags to each tier and set up firewall rules to allow the desired traffic flow

Suggested Answer: *D*
Google Cloud Platform(GCP) enforces firewall rules through rules and tags. GCP rules and tags can be defined once and used across all regions. Reference: https://cloud.google.com/docs/compare/openstack/ https://aws.amazon.com/it/blogs/aws/building-three-tier-architectures-with-security-groups/

Community vote distribution
D (100%)

shandy Highly Voted 6 years, 8 months ago
D. refer to target filtering. https://cloud.google.com/solutions/best-practices-vpc-design
upvoted 37 times

shandy Highly Voted 6 years, 8 months ago
D. refer to target filtering. https://cloud.google.com/solutions/best-practices-vpc-designD is okThank you for the link. Precisely: https://cloud.google.com/solutions/best-practices-vpc-design#target_filteringD, firewalls can be done on ip or network tags or service accounts in GCE.D is right
upvoted 37 times

tartar 6 years ago
D is ok
upvoted 8 times

pepYash 5 years, 9 months ago
Thank you for the link. Precisely: https://cloud.google.com/solutions/best-practices-vpc-design#target_filtering
upvoted 9 times

nitinz 5 years, 5 months ago
D, firewalls can be done on ip or network tags or service accounts in GCE.
upvoted 4 times

AzureDP900 3 years, 10 months ago
D is right
upvoted 1 times

amxexam 4 years, 11 months ago
Let's go with option elimination A. Add each tier to a different subnetwork >> Adding tiers to different subnets does not prevent or block them from accessing each other. Until specific firewall rules on VM or subnet allow access traffic on a specific port in the rule. B. Set up software-based firewalls on individual VMs >> Not a recommended practice will have to enable firewall anyway. C. Add tags to each tier and set up routes to allow the desired traffic flow >> Can be done but. D. Add tags to each tier and set up firewall rules to allow the desired traffic flow >> Recommended way Hence D
upvoted 11 times

Ekramy_Elnaggar Most Recent 1 year, 9 months ago
1. Firewall rules for security: Firewall rules provide the most granular and robust control over network traffic. By using tags to identify instances in each tier (web, API, database), you can create firewall rules that explicitly allow or deny traffic between these tiers. 2. Controlling traffic flow: You can create rules that: - Allow traffic from the web tier to the API tier. - Allow traffic from the API tier to the database tier. - Explicitly deny traffic between the web and database tiers. 3. Scalability and Flexibility: This approach works well even when your tiers scale independently. As new instances are added, they inherit the tags and automatically adhere to the defined firewall rules.
upvoted 1 times

ddatta 1 year, 10 months ago

D is correct

upvoted 1 times

lisabisa 2 years, 5 months ago

Routes are typically used for directing traffic between networks rather than within the same network. While tags can be used for identifying resources, they are typically used in conjunction with firewall rules for controlling traffic flow.

upvoted 2 times

AdityaGupta 2 years, 10 months ago

Why to implement anything else when Firewall is built-in within VPC and works based on Tags associated with resources.

upvoted 1 times

heretolearnazure 2 years, 12 months ago

separate vnet is ruled out as they are on same network.

upvoted 1 times

red_panda 3 years, 2 months ago

For me most suitable answer is D

upvoted 1 times

omermahgoub 3 years, 8 months ago

It's D To configure the network so that traffic flows through the web to the API tier and then on to the database tier, but does not flow between the web and the database tier, you can add tags to each tier and set up firewall rules to allow the desired traffic flow. By adding tags to each tier, you can identify the VMs that belong to each tier and create firewall rules that allow traffic between the tiers as needed. For example, you can create a firewall rule that allows traffic from the web tier to the API tier, and another rule that allows traffic from the API tier to the database tier. This will ensure that traffic flows through the desired path and is not allowed between the web and database tiers. Other options, such as adding each tier to a different subnetwork or setting up software-based firewalls on individual VMs, may not provide the necessary level of control over the traffic flow between the tiers. Setting up routes to allow the desired traffic flow may not be sufficient to prevent traffic between the web and database tiers.

upvoted 5 times

megumin 3 years, 9 months ago

D is ok

upvoted 1 times

Mahmoud_E 3 years, 10 months ago

D is right answer

upvoted 1 times

minmin2020 3 years, 10 months ago

Selected Answer: D

Having 3-tier web application deployed in the same network is wrong to begin with. However, even in different subnets you will need to apply firewall rules to prevent traffic between selected subnets. In this case they will probably be better of with D.

upvoted 1 times

minmin2020 3 years, 10 months ago

Having 3-tier web application deployed in the same network is wrong to begin with. However, even in different subnets you will need to apply firewall rules to prevent traffic between selected subnets. In this case they will probably be better of with D.Did this appear in the exam?

upvoted 1 times

zr79 3 years, 10 months ago

Did this appear in the exam?

upvoted 1 times

holerina 3 years, 11 months ago

use firewall rules

upvoted 1 times

amxexam 4 years, 3 months ago

A per my comment below .

upvoted 1 times

vincy2202 4 years, 7 months ago

D is the correct answer

upvoted 3 times

haroldbenites 4 years, 8 months ago

Go for D

upvoted 2 times

unnikrisb 4 years, 10 months ago

From Google practice exam question : D is correct because as instances scale, they will all have the same tag to identify the tier. These tags can then be leveraged in firewall rules to allow and restrict traffic as required, because tags can be used for both the target and source. https://cloud.google.com/vpc/docs/using-vpc https://cloud.google.com/vpc/docs/routes https://cloud.google.com/vpc/docs/add-remove-network-tags

upvoted 2 times

Question #34 *Topic 1*

Your development team has installed a new Linux kernel module on the batch servers in Google Compute Engine (GCE) virtual machines (VMs) to speed up the nightly batch process. Two days after the installation, 50% of the batch servers failed the nightly batch run. You want to collect details on the failure to pass back to the development team.
Which three actions should you take? (Choose three.)

A. Use Stackdriver Logging to search for the module log entries

B. Read the debug GCE Activity log using the API or Cloud Console

C. Use gcloud or Cloud Console to connect to the serial console and observe the logs

D. Identify whether a live migration event of the failed server occurred, using in the activity log

E. Adjust the Google Stackdriver timeline to match the failure time, and observe the batch server metrics

F. Export a debug VM into an image, and run the image on a local server where kernel log messages will be displayed on the native screen

Suggested Answer: *A, C, E*
Official Answer: ACE

Community vote distribution


rishab86 Highly Voted 5 years, 2 months ago
ACE A. Use Stackdriver Logging to search for the module log entries = Check logs C. Use gcloud or Cloud Console to connect to the serial console and observe the logs = Check grub messages, remember new kernel module was installed. E. Adjust the Google Stackdriver timeline to match the failure time, and observe the batch server metrics = Zoom into the time window when problem happened.
upvoted 48 times

rishab86 Highly Voted 5 years, 2 months ago
ACE A. Use Stackdriver Logging to search for the module log entries = Check logs C. Use gcloud or Cloud Console to connect to the serial console and observe the logs = Check grub messages, remember new kernel module was installed. E. Adjust the Google Stackdriver timeline to match the failure time, and observe the batch server metrics = Zoom into the time window when problem happened.But the assumption you made is that stack driver was already installed on the vms. What if it was not there? Would there be any scope to install later and retrieve the logs?A, B, E C - doesn't look correct as it ends with "observe the logs" - question is on sharing the details to development team, not to look for causeE observe the logs too. One of my problem is that several of the options go on and do some observation instead of just delivering.
upvoted 48 times

Pokchok 5 years, 2 months ago
But the assumption you made is that stack driver was already installed on the vms. What if it was not there? Would there be any scope to install later and retrieve the logs?
upvoted 3 times

AmitAr 4 years, 3 months ago
A, B, E C - doesn't look correct as it ends with "observe the logs" - question is on sharing the details to development team, not to look for cause
upvoted 1 times

AmitAr 4 years, 3 months ago
A, B, E C - doesn't look correct as it ends with "observe the logs" - question is on sharing the details to development team, not to look for causeE observe the logs too. One of my problem is that several of the options go on and do some observation instead of just delivering.
upvoted 1 times

tocsa 2 years, 2 months ago
E observe the logs too. One of my problem is that several of the options go on and do some observation instead of just delivering.
upvoted 1 times

haroldbenites 4 years, 8 months ago
Go for A,B,E. C is when the VM is running , but in this case the sentence says “recollect”. It means that “error ever” already happened.and how will activity log help?
upvoted 12 times

haroldbenites 4 years, 8 months ago

Go for A,B,E. C is when the VM is running , but in this case the sentence says “recollect”. It means that “error ever” already happened.

upvoted 12 times

pddddd 4 years, 7 months ago

and how will activity log help?

upvoted 1 times

JoseCloudEng1994 Most Recent 11 months, 3 weeks ago

B - Cannot be this as Activity Logs track GCP-level events, like instance creation, deletion, metadata changes, and API calls. They do not contain OS-level debugging info or kernel/module failure data. D - Host Errors might cause 1 crash randomly. I see this happening at work all the time. But a bunch of them at once. Unlikely F - We know that this is only happening 50% of the time. How will you reproduce this locally? The other 50% of the time it works flawlessly

upvoted 1 times

Snafkin547 1 year, 2 months ago

GCE Activity log does not provide debug logs as they lack kernel application-level details, while The serial console provides low level access to VMs boot and kernel logs. Therefore, C wins in this case.

upvoted 1 times

rajrocks171 1 year, 2 months ago

A & E - Use Cloud Logging (Previously Stackdriver logging) search and filter module log entries. Provides syslog, kernel logs, app logs etc and can be filterer by VM name, time range, severity.. C - Serial console output - Kernel panics, boot failures or critical errors related to a kernel module often print messages directly to serial console before OS might even fully functional. [Option B could be helpful but for the scenario in question I won't go to GCE activities logs(Cloud Audit Logs) skipping C (Serial Console Output). It gives the information what's happening to VM from control pane's perspective, it won't directly tell me kernel module's internal errors...]

upvoted 1 times

izekc 1 year, 5 months ago

ABE is much make sense to me

upvoted 1 times

Ishu_awsguy 1 year, 7 months ago

B vs C is the question. I vote for B , since activity logs now have been moved to audit logs and they provode system event information. C is moee real time, it will be a task to extract all the logs and put it to a file and share with dev team. Anyways audit logs is already abstracting that work for you. Hence ABE is the right answer as per me

upvoted 3 times

mahi_h 1 year, 8 months ago

I chose B over C as both options trying to read/observe logs. But C looks like reading at the runtime. Give that, it's a nightly batch process, B seems to suitable post error occurrence.

upvoted 1 times

deep316 1 year, 8 months ago

A. Use Stackdriver Logging to search for the module log entries: This will help you identify any errors or issues related to the new kernel module that were logged during the batch process. C. Use gcloud or Cloud Console to connect to the serial console and observe the logs: The serial console logs can provide detailed information about the boot process and any kernel-related messages that might indicate why the batch servers failed. D. Identify whether a live migration event of the failed server occurred, using the activity log: Live migration events can sometimes cause disruptions. Checking the activity log will help you determine if this was a factor in the failures.

upvoted 2 times

Ekramy_Elnaggar 1 year, 9 months ago

A. Use Stackdriver Logging to search for the module log entries: This is crucial. Stackdriver Logging aggregates logs from various sources, including your VMs. You can search for specific entries related to the new kernel module to identify errors, warnings, or unusual behavior that might explain the failures. C. Use gcloud or Cloud Console to connect to the serial console and observe the logs: The serial console provides access to the VM's output even if the system is unresponsive. This can be invaluable for capturing kernel panic messages, boot errors, or other critical information that might not be available through standard logging channels. E. Adjust the Google Stackdriver timeline to match the failure time, and observe the batch server metrics: Stackdriver Monitoring provides detailed performance metrics for your VMs. By aligning the timeline with the failures, you can analyze CPU usage, memory consumption, disk I/O, and network activity.

upvoted 3 times

SerGCP 1 year, 9 months ago

considering that logs from the serial console might not be useful after two days: A. Use Stackdriver Logging to search for the module log entries: This will help you identify any errors or issues logged by the new kernel module. B. Read the debug GCE Activity log using the API or Cloud Console: This can provide information on significant events like reboots or live migrations that might have affected the batch process. E. Adjust the Google

Stackdriver timeline to match the failure time, and observe the batch server metrics: This helps correlate the failure with any anomalies in the server metrics, providing insights into what might have gone wrong.

upvoted 1 times

nareshthumma 1 year, 9 months ago

Answer is ACE

upvoted 1 times

Hungdv 2 years ago

Vote ACE

upvoted 1 times

Gino17m 2 years, 3 months ago

I'm not sure but I vote for ABE

upvoted 2 times

ashishdwi007 2 years, 6 months ago

ACE makes sense. A and E dont have any doubts. Questions is that if it is B, C or F. Whats' use of serial console if team does not use it for logging especially kernal related updates. that makes sense to choose C.

upvoted 1 times

CloudDom 2 years, 9 months ago

Most automated way, with C, collecting from VMs involves a lot of manual efforts

upvoted 3 times

Question #35 *Topic 1*

Your company wants to try out the cloud with low risk. They want to archive approximately 100 TB of their log data to the cloud and test the analytics features available to them there, while also retaining that data as a long-term disaster recovery backup.
Which two steps should you take? (Choose two.)

A. Load logs into Google BigQuery

B. Load logs into Google Cloud SQL

C. Import logs into Google Stackdriver

D. Insert logs into Google Cloud Bigtable

E. Upload log files into Google Cloud Storage

Suggested Answer: *A, E*
Official Answer: AE

Community vote distribution
AE (87%) CE (13%)

rishab86 Highly Voted 5 years, 2 months ago
Answer is A as they want to load logs for analytics and E for storing data in buckets for long term.
upvoted 34 times

omermahgoub Highly Voted 3 years, 8 months ago
AE To archive approximately 100 TB of log data to the cloud and test the analytics features available while also retaining the data as a long-term disaster recovery backup, you can take the following steps: E: Upload log files into Google Cloud Storage: Google Cloud Storage is a scalable, durable, and fully-managed cloud storage service that can be used to store large amounts of data. You can upload your log files to Cloud Storage to archive them in the cloud. A: Load logs into Google BigQuery: Google BigQuery is a fully-managed, cloud-native data warehouse that can be used to analyze large amounts of data quickly and efficiently. You can load your log data into BigQuery to perform analytics on it and test the available analytics features. Other options, such as loading logs into Google Cloud SQL, importing logs into Google Stackdriver, or inserting logs into Google Cloud Bigtable, may not provide the necessary functionality for archiving and analyzing the log data.
upvoted 13 times

Ekramy_Elnaggar Most Recent 1 year, 9 months ago
A. Load logs into Google BigQuery: BigQuery is Google Cloud's serverless, highly scalable, and cost-effective multicloud data warehouse designed for data analytics. It's ideal for storing and analyzing large volumes of log data (100 TB in this case). You can use BigQuery's powerful SQL capabilities to run queries, generate reports, and gain insights from your logs. E. Upload log files into Google Cloud Storage: Cloud Storage provides durable, scalable, and secure object storage. It's perfect for storing your log data as a long-term disaster recovery backup. Cloud Storage offers different storage classes to optimize costs based on your data access frequency and retention needs.
upvoted 3 times

ionescuandrei 3 years, 4 months ago
This looks right.
upvoted 2 times

CMata 3 years, 9 months ago
If you want to analize those logs its recommended Big Query. For storing and backup Cloud Storage is your option, so AE
upvoted 3 times

AzureDP900 3 years, 10 months ago
A and E can do the required task
upvoted 2 times

minmin2020 3 years, 10 months ago
AE are the only options for analysis and archiving
upvoted 1 times

holerina 3 years, 11 months ago
A load in Big query for analytics and E for cloud storage
upvoted 1 times

Ramheadhunter 4 years ago

The key word is 'Analytics' here the main reason for moving logs to GCP is to perform Analytics on the data. BigQuery is the best suite for it. For long term storage it wiould be GC

upvoted 1 times

raaj_p 4 years ago

Selected Answer: CE

Answer is C and E

Key features Real-time log management and analysis Cloud Logging is a fully managed service that performs at scale and can ingest application and platform log data, as well as custom log data from GKE environments, VMs, and other services inside and outside of Google Cloud. Get advanced performance, troubleshooting, security, and business insights with Log Analytics, integrating the power of BigQuery into Cloud Logging. - https://cloud.google.com/products/operations

upvoted 2 times

raaj_p 4 years ago

Answer is C and E Key features Real-time log management and analysis Cloud Logging is a fully managed service that performs at scale and can ingest application and platform log data, as well as custom log data from GKE environments, VMs, and other services inside and outside of Google Cloud. Get advanced performance, troubleshooting, security, and business insights with Log Analytics, integrating the power of BigQuery into Cloud Logging. - https://cloud.google.com/products/operationsThe key word is 'Analytics' here the main reason for moving logs to GCP is to perform Analytics on the data. BigQuery is the best suite for it. For long term storage it wiould be GCSIs it even possible to import 10TB of logs into StackDriver?you don't do realtime log management on 10 TB data. You only perform analytics on it. So A for Analytics E for storage.

upvoted 2 times

Ramheadhunter 4 years ago

The key word is 'Analytics' here the main reason for moving logs to GCP is to perform Analytics on the data. BigQuery is the best suite for it. For long term storage it wiould be GCS

upvoted 1 times

tocsa 2 years, 2 months ago

Is it even possible to import 10TB of logs into StackDriver?

upvoted 1 times

AMohanty 4 years ago

you don't do realtime log management on 10 TB data. You only perform analytics on it. So A for Analytics E for storage.

upvoted 6 times

faagee01 4 years ago

Big Query for analytics, Cloud Storage for long term archive

upvoted 3 times

Dhiraj03 4 years, 2 months ago

For Storage GCS is the best option and for analyzing the data BIg Query makes sense

upvoted 1 times

Nirca 4 years, 3 months ago

A E are ok

upvoted 2 times

vincy2202 4 years, 7 months ago

AE are the correct answers

upvoted 3 times

andeu 4 years, 8 months ago

Answers: A is correct because BigQuery is the fully managed cloud data warehouse for analytics and supports the analytics requirement. E is correct because Cloud Storage provides the Coldline storage class to support long-term storage with infrequent access, which would support the long-term disaster recovery backup requirement. https://cloud.google.com/bigquery/ https://cloud.google.com/stackdriver/ https://cloud.google.com/storage/docs/storage-classes#coldline https://cloud.google.com/sql/ https://cloud.google.com/bigtable/

upvoted 4 times

andeu 4 years, 8 months ago

Answers: A is correct because BigQuery is the fully managed cloud data warehouse for analytics and supports the analytics requirement. E is correct because Cloud Storage provides the Coldline storage class to support long-term storage with infrequent access, which would support the long-term disaster recovery backup requirement. https://cloud.google.com/bigquery/ https://cloud.google.com/stackdriver/ https://cloud.google.com/storage/docs/storage-classes#coldline https://cloud.google.com/sql/ https://cloud.google.com/bigtable/But BigQuery is for SQL DATA, the logs are nosql ? Why not choose stackdriver?

upvoted 4 times

MQQ 4 years, 2 months ago

But BigQuery is for SQL DATA, the logs are nosql ? Why not choose stackdriver?

upvoted 1 times

haroldbenites 4 years, 8 months ago

Go for A,E

upvoted 3 times

Bakili 4 years, 8 months ago

A and E

upvoted 1 times

Question #36 *Topic 1*

You created a pipeline that can deploy your source code changes to your infrastructure in instance groups for self-healing. One of the changes negatively affects your key performance indicator. You are not sure how to fix it, and investigation could take up to a week.
What should you do?

A. Log in to a server, and iterate on the fox locally

B. Revert the source code change, and rerun the deployment pipeline

C. Log into the servers with the bad code change, and swap in the previous code

D. Change the instance group template to the previous one, and delete all instances

Suggested Answer: *B*
Official Answer: B

Community vote distribution

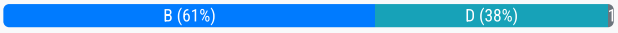

amxexam Highly Voted 4 years, 11 months ago

Let's go with option elimination A. Log in to a server, and iterate on the fix locally >> Long step, hence eliminate B. Revert the source code change and rerun the deployment pipeline >> This revert will be logged in the source repo. Will go with this way although D also is correct. C. login to the servers with the bad code change, and swap in the previous code >> C is manually doing what can be automatically done by B and C, hence eliminate. D. Change the instance group template to the previous one and delete all instances >> This is similar to B but why manually do something which is automated. Hence eliminate. But is also correct. But B is better from code lifecycle perspective. Hence BThe question itself looks the madeup. Not a real scenario ..."You created a pipeline that can deploy your source code changes to your infrastructure in instance groups for self-healing. One of the changes negatively affects your key performance indicator. " How a self healing code is affecting KPI. What was KPI, we dont know. Was the self healing done? we dont know. Dont know who make this questions. Even if we go whatever they try to ask, with options available, B is safest. However this option is just answer to any troubleshooting step. I m not convinced for the person who wrote this question

upvoted 78 times

amxexam Highly Voted 4 years, 11 months ago

Let's go with option elimination

A. Log in to a server, and iterate on the fix locally >> Long step, hence eliminate

B. Revert the source code change and rerun the deployment pipeline >> This revert will be logged in the source repo. Will go with this way although D also is correct.

C. login to the servers with the bad code change, and swap in the previous code >> C is manually doing what can be automatically done by B and C, hence eliminate.

D. Change the instance group template to the previous one and delete all instances >> This is similar to B but why manually do something which is automated. Hence eliminate. But is also correct. But B is better from code lifecycle perspective.

Hence B

upvoted 78 times

ashishdwi007 2 years, 6 months ago

The question itself looks the madeup. Not a real scenario ..."You created a pipeline that can deploy your source code changes to your infrastructure in instance groups for self-healing. One of the changes negatively affects your key performance indicator. " How a self healing code is affecting KPI. What was KPI, we dont know. Was the self healing done? we dont know. Dont know who make this questions. Even if we go whatever they try to ask, with options available, B is safest. However this option is just answer to any troubleshooting step. I m not convinced for the person who wrote this question

upvoted 3 times

ewredtrfygi 6 years ago

Too many responses saying B is the answer - I wonder if GCP pays people to provide the wrong answers on this website. It's clearly D, MIG templates support versioning, they were created to solve this exact problem. You simply select the previous template version, set that as the new deployment, and it will roll back the KPI depriving deployment and roll out the previous working deployment. The only part of D I don't like is the "terminate all instances" since you should engage in a rolling deployment, but if it's not a live website I suppose that would be fine.
https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups

upvoted 63 times

ewredtrfygi 6 years ago

Too many responses saying B is the answer - I wonder if GCP pays people to provide the wrong answers on this website. It's clearly D, MIG templates support versioning, they were created to solve this exact problem. You simply select the previous template version, set that as the new deployment, and it will roll back the KPI depriving deployment and roll out the previous working deployment. The only part of D I don't like is the "terminate all instances" since you should engage in a rolling deployment, but if it's not a live website I suppose that would be fine.
https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groupsB. keyword is "self-healing" not "auto-healing" - which means MIG not used. So correct answer is Bseems with approach, there will be a mismatch in pipelineIf you change manually the template.. why are using pipelines? B is the best answer because is automated!!! Why Google will be interested to vote the wrong answers??? They want more professionals with GCP certifications!!!!"....One of the changes has impacted negatively your PKI". Why is the question about pipeline? It is about how to do investigations and keep your PKI at the proper SLA/SLO.it mentioned about pipeline, where changes started from the code, rolling back to previous version from MIG will break the pipeline, introducing inconsistency.

upvoted 63 times

AmitAr 4 years, 2 months ago

B. keyword is "self-healing" not "auto-healing" - which means MIG not used. So correct answer is B

upvoted 2 times

Bill831231 4 years, 10 months ago

seems with approach, there will be a mismatch in pipeline

upvoted 4 times

Meyucho 4 years, 7 months ago

If you change manually the template.. why are using pipelines? B is the best answer because is automated!!! Why Google will be interested to vote the wrong answers??? They want more professionals with GCP certifications!!!!

upvoted 7 times

Meyucho 4 years, 7 months ago

If you change manually the template.. why are using pipelines? B is the best answer because is automated!!! Why Google will be interested to vote the wrong answers??? They want more professionals with GCP certifications!!!!"....One of the changes has impacted negatively your PKI". Why is the question about pipeline? It is about how to do investigations and keep your PKI at the proper SLA/SLO.

upvoted 7 times

Davidik79 4 years, 5 months ago

"....One of the changes has impacted negatively your PKI". Why is the question about pipeline? It is about how to do investigations and keep your PKI at the proper SLA/SLO.

upvoted 1 times

GreenTick 4 years, 1 month ago

it mentioned about pipeline, where changes started from the code, rolling back to previous version from MIG will break the pipeline, introducing inconsistency.

upvoted 1 times

JoseCloudEng1994 Most Recent 11 months, 3 weeks ago

A - Can be done if you are not in a rush and its not too complicated. Stepping function by function until you find the issue can take a while and if you have applied the change to 100+ servers. Good luck doing them 1 by 1 B - Right answer. Reverts your changes to prior and can be used to produce a new image with working code. C - Manual work?? If you have a bunch of servers you are liable to human error. Then you have to explain to the customer why the change only failed for them. No thanks, either apply it everywhere or you have issues D. I mean yeah but what if you dont have an instance group?

upvoted 1 times

Jorgehr7 1 year, 1 month ago

Its D. Less effort option to solve the case

upvoted 1 times

rajrocks171 1 year, 2 months ago

I would go with D. Key word "Instance groups of Self healing" - That's MIG which provides Rolling Updates option. Controlled Rollout, Zero Downtime. You just need to identify the previous stable instance template and trigger rolling updates. This will take care instance deletions after successful spin up of new instances.

upvoted 2 times

user263263 1 year, 6 months ago

If you have a pipeline that deploys, changing the deployment manually is an anti-pattern - it causes the code in the repo to be inconsistent with the code that runs. Fix forward is not an option ("investigation can take a week") so revert source code change is the only possibility.

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Quickest path to recovery: Reverting the code change that caused the performance degradation is the fastest way to restore your key performance indicator (KPI) to its previous level. Your pipeline is designed to automate deployments, so re-deploying the known good version should be straightforward. 2. Minimizes downtime: While investigating the issue is important, it can take time. Reverting first minimizes the duration of the negative impact on your KPI. 3. Clean and controlled: This approach avoids making ad-hoc changes directly on servers (options A and C), which can lead to inconsistencies and make it harder to track the problem's source. 4. Maintains instance group integrity: Option D involves deleting instances, which can disrupt services and lead to unnecessary re-creation of resources.

upvoted 2 times

sim7243 1 year, 9 months ago

option B

upvoted 1 times

dija123 2 years, 4 months ago

Agree with B

upvoted 1 times

lisabisa 2 years, 5 months ago

D is infrastructure change. B is application change. So B is correct.

upvoted 3 times

kip21 2 years, 7 months ago

D The question is talking abt MIG and you can revert Inst Template same as B. Since this is about MIG's I will choose D

upvoted 1 times

adoyt 2 years, 7 months ago

D. This is a question about instance groups and so modifying templates should be what we're looking for.

upvoted 1 times

simiramis221 2 years, 8 months ago

This a B for sure

upvoted 1 times

AdityaGupta 2 years, 10 months ago

The popped up with source code changes, hence reverting the change and deployment will solve the issue. B. Revert the source code change, and rerun the deployment pipeline

upvoted 2 times

yilexar 2 years, 10 months ago

D is incorrect: - MIG instance group template is immutable, there is no version concept. The context never mentioned that team created multiple instance group templates. - Software code change might not all ended up in the instance group templates, it depends on how the deployment pipeline is configured. Regardless, B is a best practices, ensure that your infrastructure is synced with your source control system.

upvoted 1 times

rusll 2 years, 11 months ago

You don't need to touch your code, just deploy and older version and fix the code, then deploy the fixed version

upvoted 2 times

jalberto 3 years ago

The most secure option is D Revert a source code change could be complex (if change was made in various components)

upvoted 2 times

Question #37 *Topic 1*

Your organization wants to control IAM policies for different departments independently, but centrally.
Which approach should you take?

A. Multiple Organizations with multiple Folders

B. Multiple Organizations, one for each department

C. A single Organization with Folders for each department

D. A single Organization with multiple projects, each with a central owner

Suggested Answer: *C*
Folders are nodes in the Cloud Platform Resource Hierarchy. A folder can contain projects, other folders, or a combination of both. You can use folders to group projects under an organization in a hierarchy. For example, your organization might contain multiple departments, each with its own set of GCP resources. Folders allow you to group these resources on a per-department basis. Folders are used to group resources that share common IAM policies. While a folder can contain multiple folders or resources, a given folder or resource can have exactly one parent.
Reference: https://cloud.google.com/resource-manager/docs/creating-managing-folders

Community vote distribution
C (100%)

AWS56 Highly Voted 6 years, 9 months ago
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations

I will stick with C
upvoted 27 times

AWS56 Highly Voted 6 years, 9 months ago
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations I will stick with Chttps://cloud.google.com/architecture/identity/best-practices-for-planning#use_a_separate_staging_organizationhttps://cloud.google.com/architecture/identity/best-practices-for-planning#use_organizations_to_delineate_administrative_authorityC is recommended approach
upvoted 27 times

CamiloJrJr Most Recent 1 year, 8 months ago
https://cloud.google.com/architecture/identity/best-practices-for-planning#use_a_separate_staging_organization
upvoted 1 times

CamiloJrJr 1 year, 8 months ago
https://cloud.google.com/architecture/identity/best-practices-for-planning#use_organizations_to_delineate_administrative_authority
upvoted 1 times

AzureDP900 3 years, 10 months ago
C is recommended approach
upvoted 2 times

BiddlyBdoyng 3 years, 10 months ago
The reason it isn't D is that a Dept modelled as a project puts a massive constraint on the dept that they can only have a single project, it's likely a department will want many projects.
upvoted 7 times

Ekramy_Elnaggar 1 year, 9 months ago
1. Centralized Control: A single organization provides a central point for managing all your Google Cloud resources and IAM policies. This simplifies administration and ensures consistency across your organization. 2. Independent Department Management: Folders allow you to group projects within your organization and delegate administrative control to different departments. Each department can manage its own IAM policies within its assigned folder, providing the necessary independence. 3. Hierarchical Structure: Folders provide a hierarchical structure for organizing your resources. You can create sub-folders within departments for further granularity and control. 4. Efficient Resource Management: This structure makes it easier to manage resources, track costs, and enforce security policies across your organization.
upvoted 4 times

Shasha1 1 year, 10 months ago

B

upvoted 1 times

vyomkeshbakshi 3 years, 1 month ago

Clearly C

upvoted 1 times

red_panda 3 years, 2 months ago

Is obviously C

upvoted 1 times

omermahgoub 3 years, 8 months ago

C. A single Organization with Folders for each department To control IAM policies for different departments independently but centrally, you should create a single organization and use folders to organize the policies for each department. This approach allows you to centralize the management of IAM policies for all departments within a single organization, while also allowing you to set up different policies for each department as needed. Option A, multiple organizations with multiple folders, would not be an effective solution because it would create unnecessary complexity and make it more difficult to centralize the management of IAM policies. Option B, multiple organizations, one for each department, would also not be an effective solution because it would create unnecessary complexity and make it more difficult to centralize the management of IAM policies. Option D, a single organization with multiple projects, each with a central owner, would not be an effective solution because it would not allow you to set up different policies for each department as needed.

upvoted 2 times

MarcoEscanor 3 years, 10 months ago

C. It's a best practice and I've done this with my previous and current company :)

upvoted 3 times

holerina 3 years, 11 months ago

C single org and multiple folders

upvoted 1 times

cmamiusa 4 years, 4 months ago

C is the right answer

upvoted 1 times

mygcpjourney2712 4 years, 4 months ago

Will go for C

upvoted 1 times

haroldbenites 4 years, 8 months ago

Go for C

upvoted 2 times

vincy2202 4 years, 8 months ago

C is the right answer

upvoted 1 times

nansi 4 years, 10 months ago

C shall be the correct answer

upvoted 1 times

rikoko 5 years ago

C. Seems to be best practice (cf AWS56). And I believe that D should be excluded because it says "Project owner" - it is not best practice since it's a basic role + it's not even stated as a requisite

upvoted 1 times

victory108 5 years, 3 months ago

C. A single Organization with Folders for each department

upvoted 3 times

un 5 years, 3 months ago

C is correct

upvoted 1 times

Question #38 *Topic 1*

You deploy your custom Java application to Google App Engine. It fails to deploy and gives you the following stack trace.
What should you do?

```
java.lang.SecurityException: SHA1 digest error for
com/Altostrat/CloakedServlet.class
     at com.google.appengine.runtime.Request.process
-d36f818a24b8cf1d (Request.java)
     at
sun.security.util.ManifestEntryVerifier.verify
(ManifestEntryVerifier.java:210)
     at java.util.jar.JarVerifier.processEntry
(JarVerifier.java:218)
     at java.util.jar.JarVerifier.update
(JarVerifier.java:205)
     at
java.util.jar.JarVerifiersVerifierStream.read
(JarVerifier.java:428)
     at sun.misc.Resource.getBytes
(Resource.java:124)
     at java.net.URL.ClassLoader.defineClass
(URLClassLoader.java:273)
     at sun.reflect.GeneratedMethodAccessor5.invoke
(Unknown Source)
     at
sun.reflect.DelegatingMethodAccessorImpl.invoke
(DelegatingMethodAccessorImpl.java:43)
     at java.lang.reflect.Method.invoke
(Method.java:616)
     at java.lang.ClassLoader.loadClass
(ClassLoader.java:266)
```

A. Upload missing JAR files and redeploy your application.

B. Digitally sign all of your JAR files and redeploy your application

C. Recompile the CLoakedServlet class using and MD5 hash instead of SHA1

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (75%) A (25%)

Eroc Highly Voted 6 years, 9 months ago
Signing the JAR files grants it permissions. (https://docs.oracle.com/javase/tutorial/deployment/jar/signindex.html)
upvoted 24 times

Eroc Highly Voted 6 years, 9 months ago
Signing the JAR files grants it permissions. (https://docs.oracle.com/javase/tutorial/deployment/jar/signindex.html)B is okWhere do you go? when we need you for other questions. Plz ans other q's if you have timeB, SHA1 Digest error in the first line in the error code. With Java errors, always focus on the first line in the error code, rest of the lines are garbage **mostly**.
upvoted 24 times

tartar 6 years ago
B is ok

upvoted 9 times

tartar 6 years ago

B is okWhere do you go? when we need you for other questions. Plz ans other q's if you have time

upvoted 9 times

Urban_Life 4 years, 10 months ago

Where do you go? when we need you for other questions. Plz ans other q's if you have time

upvoted 2 times

nitinz 5 years, 5 months ago

B, SHA1 Digest error in the first line in the error code. With Java errors, always focus on the first line in the error code, rest of the lines are garbage **mostly**.

upvoted 19 times

omermahgoub 3 years, 8 months ago

The most likely cause of the error is that one of the JAR files in your application has been tampered with or is corrupt. The SHA1 digest error indicates that the JAR file's signature does not match the expected value, which could be due to tampering or corruption. To fix the issue, you should try uploading missing JAR files and redeploying your application. If the issue persists, you may need to digitally sign all of your JAR files and redeploy your application to ensure that the signatures are valid. You should not try to recompile the Cloaked

upvoted 15 times

Ekramy_Elnaggar Most Recent 1 year, 9 months ago

1. JAR signing and integrity: Digitally signing your JAR files ensures their authenticity and integrity. It adds a digital signature that verifies the origin and confirms that the file hasn't been tampered with. This is crucial for security and prevents issues like the SHA1 digest error you're encountering. 2. App Engine requirement: Google App Engine enforces JAR signing for security reasons. All deployed applications must have properly signed JAR files.

upvoted 5 times

chrissamharris 1 year, 9 months ago

A SHA1 digest error during the deployment of a Java application to Google App Engine (GAE) typically indicates an issue with the integrity of your JAR files. This error can arise due to corrupted, modified, or missing JAR files that are essential for your application to run correctly.

upvoted 1 times

lisabisa 2 years, 5 months ago

A missing JAR (Java ARchive) file indicates a problem with the code to be deployed. B Digitally signing all of your JAR files indicates a problem with the signature. A is better

upvoted 3 times

AdityaGupta 2 years, 10 months ago

"JavaVerifier.Java.428" is the key here

upvoted 2 times

SantoshPanigrahi 2 years, 12 months ago

B is the correct answer.

upvoted 1 times

JC0926 3 years, 5 months ago

Selected Answer: A

A. Upload missing JAR files and redeploy your application.

The error message indicates that there is a problem with the SHA1 digest for the "com/altostrat/cloakedservlet.class" file. This can be caused by a corrupted or incomplete JAR file. Therefore, the best course of action is to upload any missing JAR files and redeploy the application.

upvoted 2 times

JC0926 3 years, 5 months ago

A. Upload missing JAR files and redeploy your application. The error message indicates that there is a problem with the SHA1 digest for the "com/altostrat/cloakedservlet.class" file. This can be caused by a corrupted or incomplete JAR file. Therefore, the best course of action is to upload any missing JAR files and redeploy the application.Corrupted or incomplete. But not missing.

upvoted 2 times

tocsa 2 years, 2 months ago

Corrupted or incomplete. But not missing.

upvoted 2 times

Racinely 3 years, 9 months ago

B is correct

upvoted 1 times

minmin2020 3 years, 10 months ago

Selected Answer: B

Ok B but how is this question related to a GCP exam? I guess a google search will be faster than reading the theory around Java (unless you are a developer).

upvoted 7 times

minmin2020 3 years, 10 months ago

Ok B but how is this question related to a GCP exam? I guess a google search will be faster than reading the theory around Java (unless you are a developer).have you done any Azure exams? you will thank Googlenothing to do with GCP however basic troubleshooting skills required as a DevOps or Architect, B is fineI see your point, but for basic troubleshooting of apps, i will usually have access to google (aka stackoverflow homepage). This could have been a cloud developer question that they repurposed.

upvoted 7 times

zr79 3 years, 10 months ago

have you done any Azure exams? you will thank Google

upvoted 5 times

AzureDP900 3 years, 10 months ago

nothing to do with GCP however basic troubleshooting skills required as a DevOps or Architect, B is fine

upvoted 3 times

AzureDP900 3 years, 10 months ago

nothing to do with GCP however basic troubleshooting skills required as a DevOps or Architect, B is fineI see your point, but for basic troubleshooting of apps, i will usually have access to google (aka stackoverflow homepage). This could have been a cloud developer question that they repurposed.

upvoted 3 times

[Removed] 3 years, 7 months ago

I see your point, but for basic troubleshooting of apps, i will usually have access to google (aka stackoverflow homepage). This could have been a cloud developer question that they repurposed.

upvoted 1 times

holerina 3 years, 11 months ago

B is the right answer

upvoted 1 times

avinashvidyarthi 4 years, 3 months ago

B is Correct!

upvoted 1 times

Munna19 4 years, 3 months ago

B is the right answer

upvoted 1 times

vincy2202 4 years, 7 months ago

B is the correct answer

upvoted 2 times

duocnh 4 years, 8 months ago

vote B

upvoted 1 times

TheCloudBoy77 4 years, 9 months ago

B is correct answer

upvoted 1 times

bala786 5 years, 1 month ago

Option B. Digitally sign all of your JAR files and redeploy your application

upvoted 1 times

Question #39 *Topic 1*

You are designing a mobile chat application. You want to ensure people cannot spoof chat messages, by providing a message were sent by a specific user.
What should you do?

A. Tag messages client side with the originating user identifier and the destination user.

B. Encrypt the message client side using block-based encryption with a shared key.

C. Use public key infrastructure (PKI) to encrypt the message client side using the originating user's private key.

D. Use a trusted certificate authority to enable SSL connectivity between the client application and the server.

Suggested Answer: *C*
Official Answer: C

Community vote distribution
C (78%) D (22%)

KouShikyou Highly Voted 6 years, 10 months ago
I am not sure about this one. D works if SSL client authentication is enabled. C works as well if client encrypts message with private key and server decrypt with public key. I prefer C.Agree with CI agree with C on this one.
upvoted 38 times

KouShikyou Highly Voted 6 years, 10 months ago
I am not sure about this one. D works if SSL client authentication is enabled. C works as well if client encrypts message with private key and server decrypt with public key. I prefer C.
upvoted 38 times

JoeShmoe 6 years, 9 months ago
Agree with C
upvoted 5 times

asfar 6 years, 7 months ago
I agree with C on this one.
upvoted 5 times

Tobbe 5 years, 6 months ago
Encrypting each block and tagging each message at the client side is an overhead on the application. Best method which has been adopted since years is contacting SSL provider and use public certificate to encrypt the traffic between client and server. D is correctIf you use server public key you aren't meeting the goal. Don't miss the "for specific user" in the statementIf you use the server's public certificate to encrypt your data you only ensure the right server is the only one to read you. But anyone can use the same encryption key as you did and pretend to be you. Hence it does not solve our authentication problematicthanks for your insight! C is correct.Can you explain why you think this?I cannot agree with you. Before one be able to pretend to be someone else, he should know his (someone's) password on the Chat Server...SSL doesn't use server's public key to encrypt data. This is definitely wrong. Please read SSL specs. SSL uses a separate session key for message encryption. This session key is temporary and will be rotated for every single session.
upvoted 14 times

Tobbe 5 years, 6 months ago
Encrypting each block and tagging each message at the client side is an overhead on the application. Best method which has been adopted since years is contacting SSL provider and use public certificate to encrypt the traffic between client and server.

D is correct
upvoted 14 times

Meyucho 4 years, 7 months ago
If you use server public key you aren't meeting the goal. Don't miss the "for specific user" in the statement
upvoted 1 times

Alekshar 5 years, 5 months ago
If you use the server's public certificate to encrypt your data you only ensure the right server is the only one to read you. But anyone can use the same encryption key as you did and pretend to be you. Hence it does not solve our authentication problematicthanks for your insight! C is correct.Can you

explain why you think this?I cannot agree with you. Before one be able to pretend to be someone else, he should know his (someone's) password on the Chat Server...SSL doesn't use server's public key to encrypt data. This is definitely wrong. Please read SSL specs. SSL uses a separate session key for message encryption. This session key is temporary and will be rotated for every single session.

upvoted 6 times

Alekshar 5 years, 5 months ago

If you use the server's public certificate to encrypt your data you only ensure the right server is the only one to read you.

But anyone can use the same encryption key as you did and pretend to be you. Hence it does not solve our authentication problematic

upvoted 6 times

Tobbe 5 years, 5 months ago

thanks for your insight! C is correct.Can you explain why you think this?

upvoted 2 times

Tobbe 5 years, 5 months ago

thanks for your insight! C is correct.

upvoted 2 times

stefanop 2 years, 8 months ago

Can you explain why you think this?

upvoted 1 times

lynx256 5 years, 4 months ago

I cannot agree with you. Before one be able to pretend to be someone else, he should know his (someone's) password on the Chat Server...

upvoted 1 times

PeppaPig 5 years ago

SSL doesn't use server's public key to encrypt data. This is definitely wrong. Please read SSL specs. SSL uses a separate session key for message encryption. This session key is temporary and will be rotated for every single session.

upvoted 4 times

hpf97 Most Recent 1 year, 6 months ago

A : a tag can be spoofed by anybody B : so for a shared key D : can only secure the communication but ensure the sender's identity C : good answer

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Digital Signatures and Non-Repudiation: PKI provides the foundation for digital signatures. When a user sends a message, it's encrypted with their private key. The recipient can then use the sender's public key to decrypt it. This ensures: - Authenticity: The message truly originated from the claimed sender. - Non-repudiation: The sender cannot deny sending the message. - Integrity: The message hasn't been tampered with in transit. 2. How it prevents spoofing: Since only the sender has access to their private key, no one else can create a message that would decrypt correctly with their public key. This effectively prevents spoofing.

upvoted 3 times

Chojrak 2 years, 2 months ago

The "C" answer is either messed up on purpose, or somebody dumped it wrong. When you use PKI (Public Key Infrastructure), you encrypt using a _public_ key of the recipient, and the recipient decrypts using their _private_ key. Sample reference that this is correct: https://www.keyfactor.com/education-center/what-is-pki/ On the contrary, when a messages is _digitally signed_, the originator is using their _private_ key to sign the message, and the recipient is verifying it using _public key_ of the _originator_. I still don't know which answer would I choose on the actual exam.

upvoted 4 times

yas_cloud 2 years, 6 months ago

Option A is not secure because anyone who intercepts the message could modify the user identifiers. Option B does not provide a way to verify the sender's identity. Option D is important for securing the connection between the client and server, but it does not prevent message spoofing by itself. Hence C.

upvoted 1 times

02fc23a 2 years, 8 months ago

I was considering D, but nope, C: https://support.google.com/messages/answer/10262381#:~:text=your%20device%20and%20the%20device%20you%20message

upvoted 1 times

_kartik_raj 2 years, 9 months ago

Answer : D Let me clarify , what PKI is saying i think first the answer is D , reason , Just understand what it says , i.e. option c - Using PKI to encrypt messages using the originating user's private key, now couple people are saying that it is good and then the server will decrypt the msg using public key, but can't you see anyone in the whole world will be able to see the messages as public key is available publicly. ideally what should have been the

solution, Using the public key of receiver the messages should have been encrypted then the receiver would have decrypted using his private key, which absolutely makes sense, Talking about ssl i think its one of the widely used secure tech for communication between client and server

upvoted 1 times

_kartik_raj 2 years, 9 months ago

Answer : D Let me clarify , what PKI is saying i think first the answer is D , reason , Just understand what it says , i.e. option c - Using PKI to encrypt messages using the originating user's private key, now couple people are saying that it is good and then the server will decrypt the msg using public key, but can't you see anyone in the whole world will be able to see the messages as public key is available publicly. ideally what should have been the solution, Using the public key of receiver the messages should have been encrypted then the receiver would have decrypted using his private key, which absolutely makes sense, Talking about ssl i think its one of the widely used secure tech for communication between client and serverThanks for your explanation, so I choose C.....Why? Can you explain that?

upvoted 1 times

ArtistS 2 years, 9 months ago

Thanks for your explanation, so I choose C.....Why? Can you explain that?

upvoted 1 times

ArtistS 2 years, 9 months ago

Thanks for your explanation, so I choose C.....

upvoted 1 times

stefanop 2 years, 8 months ago

Why? Can you explain that?

upvoted 1 times

AdityaGupta 2 years, 10 months ago

Option C - Use public key infrastructure (PKI) to encrypt the message client-side using the originating user's private key: Using PKI to encrypt messages using the originating user's private key provides end-to-end encryption, which means only the intended recipient can decrypt the message. This option also ensures that the message's authenticity is protected. If a malicious user changes the sender's name, the recipient will not be able to decrypt the message since it was not encrypted using the correct private key. This option is a strong method for securing chat messages.

upvoted 4 times

patricklin1105 3 years ago

C is wrong because private can only be used for signing, not encrypting. Public key is used for encrypting, private key is used for decrypting.But still, spoofing is not about reading data in clear but impersonating someone else. Using D you can sign the message by confirming your identity.

upvoted 3 times

patricklin1105 3 years ago

C is wrong because private can only be used for signing, not encrypting. Public key is used for encrypting, private key is used for decrypting.

upvoted 3 times

stefanop 2 years, 8 months ago

But still, spoofing is not about reading data in clear but impersonating someone else. Using D you can sign the message by confirming your identity.

upvoted 1 times

BigfootPanda 3 years, 1 month ago

Selected Answer: C

As question is about ensuring a specific user sent a message, answer could not be D, which would ensure secure message transmission, but not message origin (which can only be done by using asymmetric key)

upvoted 3 times

BigfootPanda 3 years, 1 month ago

As question is about ensuring a specific user sent a message, answer could not be D, which would ensure secure message transmission, but not message origin (which can only be done by using asymmetric key)So the sender is signing the message while encrypting it with the private key?

upvoted 3 times

stefanop 2 years, 8 months ago

So the sender is signing the message while encrypting it with the private key?

upvoted 1 times

oriori123123 3 years, 2 months ago

bard say D. and ChatGTP say C..

upvoted 1 times

ionescuandrei 3 years, 4 months ago

I noticed this question in other tests and the suggested answer was C.

upvoted 1 times

JC0926 3 years, 5 months ago

To prevent message spoofing, it is important to ensure that messages cannot be altered or forged by anyone other than the originating user. One way to accomplish this is by using public key infrastructure (PKI) to encrypt messages using the originating user's private key.

upvoted 4 times

omermahgoub 3 years, 8 months ago

To ensure that chat messages cannot be spoofed and that the messages are truly sent by a specific user, the best option would be to use public key infrastructure (PKI) to encrypt the message client side using the originating user's private key. This would allow the recipient to verify the authenticity of the message by using the originating user's public key to decrypt the message.

Option A, tagging the message with the originating user identifier and the destination user, would not ensure the authenticity of the message, as it could potentially be forged by an attacker.

upvoted 3 times

omermahgoub 3 years, 8 months ago

To ensure that chat messages cannot be spoofed and that the messages are truly sent by a specific user, the best option would be to use public key infrastructure (PKI) to encrypt the message client side using the originating user's private key. This would allow the recipient to verify the authenticity of the message by using the originating user's public key to decrypt the message. Option A, tagging the message with the originating user identifier and the destination user, would not ensure the authenticity of the message, as it could potentially be forged by an attacker.Option B, encrypting the message using block-based encryption with a shared key, would also not ensure the authenticity of the message, as the shared key could potentially be compromised by an attacker. Option D, using a trusted certificate authority to enable SSL connectivity between the client application and the server, would help to secure the communication channel between the client and the server, but it would not necessarily ensure the authenticity of the chat messages themselves. Overall, using PKI and the originating user's private key to encrypt the message would be the most effective way to ensure the authenticity of the chat messages in your mobile chat application.

upvoted 3 times

omermahgoub 3 years, 8 months ago

Option B, encrypting the message using block-based encryption with a shared key, would also not ensure the authenticity of the message, as the shared key could potentially be compromised by an attacker. Option D, using a trusted certificate authority to enable SSL connectivity between the client application and the server, would help to secure the communication channel between the client and the server, but it would not necessarily ensure the authenticity of the chat messages themselves. Overall, using PKI and the originating user's private key to encrypt the message would be the most effective way to ensure the authenticity of the chat messages in your mobile chat application.

upvoted 2 times

FateSpringfield 3 years, 8 months ago

C is the answer, The requirement is the integrity of messages sent in CIA security (Confidentiality, Integrity, and Availability). For Confidentiality, using PublicKey of receiver, for Integrity, using PrivateKey of sender. D works in case of SSL client authentication.

upvoted 3 times

megumin 3 years, 9 months ago

C is ok

upvoted 1 times

Question #40 *Topic 1*

As part of implementing their disaster recovery plan, your company is trying to replicate their production MySQL database from their private data center to their
GCP project using a Google Cloud VPN connection. They are experiencing latency issues and a small amount of packet loss that is disrupting the replication.
What should they do?

A. Configure their replication to use UDP.

B. Configure a Google Cloud Dedicated Interconnect.

C. Restore their database daily using Google Cloud SQL.

D. Add additional VPN connections and load balance them.

E. Send the replicated transaction to Google Cloud Pub/Sub.

Suggested Answer: *B*
Official Answer: B

Community vote distribution


mawsman Highly Voted 6 years, 6 months ago
It's latency issues. That won't be solved by adding another VPN tunnel. If it was just a throughput issue then VPN would do, however to improve latency you need to go layer 2. Answer is B
upvoted 36 times

chiar Highly Voted 6 years, 9 months ago
I think B is correct. I think it is more reliable.
upvoted 30 times

desertlotus1211 Most Recent 1 year, 8 months ago
Adding VPN connections may improve bandwidth but does not resolve latency or packet loss issues caused by public internet routing... though not mentioned, we must 'think' beyond the scope of the question and ask 'what is causing the latency'... Answer B.
upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago
1. Dedicated Interconnect for high performance: Dedicated Interconnect provides a direct physical connection between your on-premises network and Google Cloud. This offers significantly lower latency, higher bandwidth, and greater reliability compared to a VPN. It's ideal for demanding workloads like database replication. 2. Reduced latency and packet loss: By bypassing the public internet, Dedicated Interconnect minimizes latency and packet loss, ensuring consistent and efficient data transfer for your MySQL replication. 3. Enhanced reliability: Dedicated Interconnect provides a more stable and predictable connection compared to a VPN, which can be affected by internet traffic fluctuations. Why D is not correct: Add additional VPN connections and load balance them: This might improve bandwidth slightly but won't address the fundamental latency and packet loss issues inherent with VPNs over the public internet.
upvoted 3 times

19040e5 2 years, 3 months ago
It's B, Interconnect. D is wrong in this case: While this might help distribute traffic, it won't solve the underlying issue of latency and packet loss caused by the inherent limitations of VPNs.
upvoted 2 times

AWS_Sam 2 years, 7 months ago
Dedicated interconnect is the answer. A second VPN will give you an HA solution, not going to resolve the latency.
upvoted 2 times

[Removed] 2 years, 7 months ago
Have you mind the budget you'll need to improve network infrastructure - whether it's dedicated interconnect or duplicating VPN connection? Mega IT corps may can afford, but I won't approve the work just for disaster recovery plan, if I were the authority. It's definatedly overkill. Main problem here is a latency issue and/or packet loss, yet the reason hasn't clearly configured. Whether it's occational, or repeatetive, and/or by DB engine or by network, mostly unknown. But you don't have to solve the problem if there's better bypass. Simply retry and test it. There C also can be a solution (which likely

being placed already), but it doesn't have significant feature for logging a transaction success/fail. If you take advance Pub/Sub, you can track each transaction processes. Ordering, can adding another issue, but it worth a try - cost effectively.

upvoted 2 times

Roro_Brother 2 years, 8 months ago

Correct answer is B as its a latency issue.

upvoted 2 times

stefanop 2 years, 8 months ago

I go for B. Latency won't be solved by adding new VPN tunnels.

upvoted 2 times

Demo_Helloworld 2 years, 11 months ago

We have something called as HAVPN which uses 2 VPN Connections at a time. As this Question is old. we dont have this Option called use HAVPN. Now its Updated so the answer will be D. As its just a replication for disaster recovery and they are facing very minimal challenges

upvoted 6 times

SandipGhosal 3 years, 5 months ago

I think the best option is E, using PubSub. In question the main issue is " a small amount of packet loss". As per google PubSub documentation, data replication among databases is one of the common use case of PubSub. The asynchronously communication of PubSub can overcome small latency issues. Setting up dedicated interconnect would be very costly and required many pre-requisites.

upvoted 3 times

Shawnn 3 years, 5 months ago

You could technically use your private key to encrypt a message, but it would not be secure because anyone who has your public key could decrypt the message. The recommended practice is to use your private key only for decryption and to use the recipient's public key for encryption. I vote for D

upvoted 1 times

SirajShan 3 years, 5 months ago

Configuring a Google Cloud Dedicated Interconnect requirement for company is a proximity to colocation facility and meeting condition to have dedicated interconnect. Had this was possible why did they used Cloud VPN in the first place ? I think answer should be D.

upvoted 1 times

n_nana 3 years, 7 months ago

If i face this issue, I will give a try with additional VPN, decision using VPN, maybe because they need encryption as well. With switching to dedicated interconnect, you have to implement your own VPN solution or application encryption. so it need more anaylsis to just skip VPN and use dedicated interconnect solution.

upvoted 2 times

omermahgoub 3 years, 8 months ago

The company should consider configuring a Google Cloud Dedicated Interconnect. A Google Cloud Dedicated Interconnect provides a private connection between the company's on-premises data center and GCP, which can help to reduce latency and improve the reliability of the connection. This can be particularly useful for replicating large amounts of data or for applications that require low-latency connectivity.

Option A, configuring the replication to use UDP, would not necessarily improve the reliability of the connection, as UDP is a connectionless protocol that does not guarantee delivery of packets.

Option C, restoring the database daily using Google Cloud SQL, would not address the underlying issues with the replication process.

upvoted 2 times

omermahgoub 3 years, 8 months ago

The company should consider configuring a Google Cloud Dedicated Interconnect. A Google Cloud Dedicated Interconnect provides a private connection between the company's on-premises data center and GCP, which can help to reduce latency and improve the reliability of the connection. This can be particularly useful for replicating large amounts of data or for applications that require low-latency connectivity. Option A, configuring the replication to use UDP, would not necessarily improve the reliability of the connection, as UDP is a connectionless protocol that does not guarantee delivery of packets. Option C, restoring the database daily using Google Cloud SQL, would not address the underlying issues with the replication process.Option D, adding additional VPN connections and load balancing them, may help to improve the reliability of the connection by providing redundancy, but it may not necessarily address latency issues. Option E, sending the replicated transaction to Google Cloud Pub/Sub, could potentially help to improve the reliability of the replication process by allowing the company to handle failures and retries in a more structured way, but it would not necessarily address latency issues. Overall, configuring a Google Cloud Dedicated Interconnect is likely to be the most effective solution for addressing latency issues and packet loss in the replication process.Why focus on latency when the solution is for disaster recovery? The main issue is packet loss. While this can be solved with Dedicated Interconnect or Cloud Pub/Sub, PubSub seems like a cheaper alternative that prevents data loss and achieves reliability. I wouldn't care about latency as the backup is only for DA.. so how does it matter if it goes slowly?Because

latency and packet loss are probably coming from traffic going over public internet. This is the Cloud VPN definition from the public documentation: "Cloud VPN securely extends your peer network to Google's network through an IPsec VPN tun

upvoted 2 times

omermahgoub 3 years, 8 months ago

Option D, adding additional VPN connections and load balancing them, may help to improve the reliability of the connection by providing redundancy, but it may not necessarily address latency issues. Option E, sending the replicated transaction to Google Cloud Pub/Sub, could potentially help to improve the reliability of the replication process by allowing the company to handle failures and retries in a more structured way, but it would not necessarily address latency issues. Overall, configuring a Google Cloud Dedicated Interconnect is likely to be the most effective solution for addressing latency issues and packet loss in the replication process.Why focus on latency when the solution is for disaster recovery? The main issue is packet loss. While this can be solved with Dedicated Interconnect or Cloud Pub/Sub, PubSub seems like a cheaper alternative that prevents data loss and achieves reliability. I wouldn't care about latency as the backup is only for DA.. so how does it matter if it goes slowly?Because latency and packet loss are probably coming from traffic going over public internet. This is the Cloud VPN definition from the public documentation: "Cloud VPN securely extends your peer network to Google's network through an IPsec VPN tunnel. Traffic is encrypted and travels between the two networks over the public internet. Cloud VPN is useful for low-volume data connections." There are documents showing what happens when public internet is used, but basically there's no way to prevent information from going through multiple hops over the internet, which is why Dedicated Interconnect should work. Plus, VPN is recommended for low-volume data connections, and I highly doubt that replicating a database is considered a low-volume operation. I'm going with B: using Cloud Dedicated Interconnect.

upvoted 1 times

omermahgoub 3 years, 8 months ago

Option D, adding additional VPN connections and load balancing them, may help to improve the reliability of the connection by providing redundancy, but it may not necessarily address latency issues.

Option E, sending the replicated transaction to Google Cloud Pub/Sub, could potentially help to improve the reliability of the replication process by allowing the company to handle failures and retries in a more structured way, but it would not necessarily address latency issues.

Overall, configuring a Google Cloud Dedicated Interconnect is likely to be the most effective solution for addressing latency issues and packet loss in the replication process.

upvoted 1 times

VSMu 3 years, 6 months ago

Why focus on latency when the solution is for disaster recovery? The main issue is packet loss. While this can be solved with Dedicated Interconnect or Cloud Pub/Sub, PubSub seems like a cheaper alternative that prevents data loss and achieves reliability. I wouldn't care about latency as the backup is only for DA.. so how does it matter if it goes slowly?Because latency and packet loss are probably coming from traffic going over public internet. This is the Cloud VPN definition from the public documentation: "Cloud VPN securely extends your peer network to Google's network through an IPsec VPN tunnel. Traffic is encrypted and travels between the two networks over the public internet. Cloud VPN is useful for low-volume data connections." There are documents showing what happens when public internet is used, but basically there's no way to prevent information from going through multiple hops over the internet, which is why Dedicated Interconnect should work. Plus, VPN is recommended for low-volume data connections, and I highly doubt that replicating a database is considered a low-volume operation. I'm going with B: using Cloud Dedicated Interconnect.

upvoted 3 times

VSMu 3 years, 6 months ago

Why focus on latency when the solution is for disaster recovery? The main issue is packet loss. While this can be solved with Dedicated Interconnect or Cloud Pub/Sub, PubSub seems like a cheaper alternative that prevents data loss and achieves reliability. I wouldn't care about latency as the backup is only for DA.. so how does it matter if it goes slowly?

upvoted 3 times

jrisl1991 2 years, 10 months ago

Because latency and packet loss are probably coming from traffic going over public internet. This is the Cloud VPN definition from the public documentation: "Cloud VPN securely extends your peer network to Google's network through an IPsec VPN tunnel. Traffic is encrypted and travels between the two networks over the public internet. Cloud VPN is useful for low-volume data connections." There are documents showing what happens when public internet is used, but basically there's no way to prevent information from going through multiple hops over the internet, which is why Dedicated Interconnect should work. Plus, VPN is recommended for low-volume data connections, and I highly doubt that replicating a database is considered a low-volume operation. I'm going with B: using Cloud Dedicated Interconnect.

upvoted 2 times

ashrafh 3 years, 9 months ago

so just to solve this issue we are going over a Dedicated Interconnect imagine saying this to a your project head.

upvoted 4 times

megumin 3 years, 9 months ago

ok for B

upvoted 2 times

Question #41 *Topic 1*

Your customer support tool logs all email and chat conversations to Cloud Bigtable for retention and analysis. What is the recommended approach for sanitizing this data of personally identifiable information or payment card information before initial storage?

A. Hash all data using SHA256

B. Encrypt all data using elliptic curve cryptography

C. De-identify the data with the Cloud Data Loss Prevention API

D. Use regular expressions to find and redact phone numbers, email addresses, and credit card numbers

Suggested Answer: *C*

Reference: https://cloud.google.com/solutions/pci-dss-compliance-in-gcp#using_data_loss_prevention_api_to_sanitize_data

Community vote distribution

C (100%)

AWS56 Highly Voted 6 years, 7 months ago

C is the answerC is okC, data sanitization = DLP

upvoted 23 times

AWS56 Highly Voted 6 years, 7 months ago

C is the answer

upvoted 23 times

tartar 6 years ago

C is ok

upvoted 8 times

nitinz 5 years, 5 months ago

C, data sanitization = DLP

upvoted 8 times

omermahgoub 3 years, 8 months ago

The recommended approach for sanitizing data of personally identifiable information or payment card information before storing it in Cloud Bigtable is option C: De-identify the data with the Cloud Data Loss Prevention API. The Cloud Data Loss Prevention (DLP) API is a powerful tool that allows you to automatically discover, classify, and redact sensitive data in your organization. It uses advanced machine learning techniques to accurately identify and protect a wide range of sensitive data types, including personal information such as names, addresses, phone numbers, and payment card information. Using the DLP API to de-identify your data before storing it in Cloud Bigtable is the most effective way to ensure that sensitive information is protected and not accessible to unauthorized users.Option A: Hashing data using SHA256 is not sufficient for protecting sensitive information, as hashes can be reversed using various techniques. Option B: Encrypting data using elliptic curve cryptography is a good option for protecting data, but it requires that you have a secure way to store and manage the encryption keys. If the keys are lost or compromised, the data will be inaccessible. Option D: Using regular expressions to find and redact phone numbers, email addresses, and credit card numbers can be effective in some cases, but it requires that you have a complete and up-to-date list of all the data patterns that you want to protect. It is also prone to errors and may not be able to detect all instances of sensitive data.About D, usually is recommended that you don´t reinvent the wheel specially when talking about security .

upvoted 16 times

omermahgoub 3 years, 8 months ago

The recommended approach for sanitizing data of personally identifiable information or payment card information before storing it in Cloud Bigtable is option C: De-identify the data with the Cloud Data Loss Prevention API.

The Cloud Data Loss Prevention (DLP) API is a powerful tool that allows you to automatically discover, classify, and redact sensitive data in your organization. It uses advanced machine learning techniques to accurately identify and protect a wide range of sensitive data types, including personal information such as names, addresses, phone numbers, and payment card information.

Using the DLP API to de-identify your data before storing it in Cloud Bigtable is the most effective way to ensure that sensitive information is protected and not accessible to unauthorized users.

upvoted 16 times

omermahgoub 3 years, 8 months ago

Option A: Hashing data using SHA256 is not sufficient for protecting sensitive information, as hashes can be reversed using various techniques.

Option B: Encrypting data using elliptic curve cryptography is a good option for protecting data, but it requires that you have a secure way to store and manage the encryption keys. If the keys are lost or compromised, the data will be inaccessible.

Option D: Using regular expressions to find and redact phone numbers, email addresses, and credit card numbers can be effective in some cases, but it requires that you have a complete and up-to-date list of all the data patterns that you want to protect. It is also prone to errors and may not be able to detect all instances of sensitive data.

upvoted 3 times

omermahgoub 3 years, 8 months ago

Option A: Hashing data using SHA256 is not sufficient for protecting sensitive information, as hashes can be reversed using various techniques. Option B: Encrypting data using elliptic curve cryptography is a good option for protecting data, but it requires that you have a secure way to store and manage the encryption keys. If the keys are lost or compromised, the data will be inaccessible. Option D: Using regular expressions to find and redact phone numbers, email addresses, and credit card numbers can be effective in some cases, but it requires that you have a complete and up-to-date list of all the data patterns that you want to protect. It is also prone to errors and may not be able to detect all instances of sensitive data.About D, usually is recommended that you don´t reinvent the wheel specially when talking about security .

upvoted 3 times

Kiroo 3 years, 3 months ago

About D, usually is recommended that you don´t reinvent the wheel specially when talking about security .

upvoted 1 times

Ekramy_Elnaggar Most Recent 1 year, 9 months ago

1. Accurate and comprehensive: The Cloud DLP API is specifically designed to identify and redact sensitive information like PII (personally identifiable information) and payment card data. It uses advanced techniques like machine learning to accurately detect sensitive data, even in complex and unstructured text. 2. Context-aware: DLP understands the context of the data. It can differentiate between a credit card number and a similar-looking sequence of numbers that's not a credit card. This reduces false positives and ensures accurate sanitization. 3. Flexible and customizable: You can configure DLP to detect specific types of sensitive data, define your own detection rules, and choose how to de-identify the data (e.g., redaction, masking, tokenization). 4. Scalable and efficient: DLP can handle large volumes of data and integrates seamlessly with other GCP services like Cloud Storage and BigQuery.

upvoted 3 times

sim7243 1 year, 9 months ago

C option

upvoted 1 times

Palan 2 years, 11 months ago

Without any doubt C is the right answer as the DLP API is a flexible and robust tool that helps identify sensitive data like credit card numbers, social security numbers, names and other forms of personally identifiable information (PII).

upvoted 1 times

shutupbot 3 years, 4 months ago

Cloud Data Loss Prevention API provides obfuscation and de-identification methods like masking and tokenization. Especially for credit card transactions, the card numbers are supposed to be tokenized. Therefore, this API is helpful.

upvoted 1 times

mbrochard 3 years, 9 months ago

C for sure !

upvoted 1 times

AniketD 3 years, 9 months ago

C is correct, DLP is the solution

upvoted 1 times

megumin 3 years, 9 months ago

ok for C

upvoted 1 times

vincy2202 4 years, 7 months ago

C is the correct answer

upvoted 2 times

haroldbenites 4 years, 8 months ago

Go for Chttps://cloud.google.com/dlp

upvoted 3 times

haroldbenites 4 years, 8 months ago

Go for C

upvoted 3 times

haroldbenites 4 years, 8 months ago

https://cloud.google.com/dlp

upvoted 1 times

MamthaSJ 5 years, 1 month ago

Answer is C

upvoted 1 times

victory108 5 years, 3 months ago

C. De-identify the data with the Cloud Data Loss Prevention API

upvoted 3 times

un 5 years, 3 months ago

C is correct

upvoted 1 times

sidhappy 5 years, 4 months ago

Effectively reduce data risk with de-identification methods like masking and tokenization https://cloud.google.com/dlp

upvoted 3 times

Ausias18 5 years, 4 months ago

Answer is C

upvoted 1 times

lynx256 5 years, 4 months ago

C is ok

upvoted 1 times

Question #42 *Topic 1*

You are using Cloud Shell and need to install a custom utility for use in a few weeks. Where can you store the file so it is in the default execution path and persists across sessions?

A. ~/bin

B. Cloud Storage

C. /google/scripts

D. /usr/local/bin

Suggested Answer: *A*

Official Answer: A

Community vote distribution

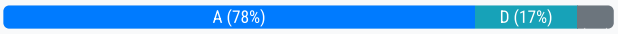

ffk Highly Voted 6 years, 9 months ago

A is correct https://cloud.google.com/shell/docs/how-cloud-shell-works Cloud Shell provisions 5 GB of free persistent disk storage mounted as your $HOME directory on the virtual machine instance. This storage is on a per-user basis and is available across projects. Unlike the instance itself, this storage does not time out on inactivity. All files you store in your home directory, including installed software, scripts and user configuration files like .bashrc and .vimrc, persist between sessions. Your $HOME directory is private to you and cannot be accessed by other users.

upvoted 77 times

ffk Highly Voted 6 years, 9 months ago

A is correct https://cloud.google.com/shell/docs/how-cloud-shell-works Cloud Shell provisions 5 GB of free persistent disk storage mounted as your $HOME directory on the virtual machine instance. This storage is on a per-user basis and is available across projects. Unlike the instance itself, this storage does not time out on inactivity. All files you store in your home directory, including installed software, scripts and user configuration files like .bashrc and .vimrc, persist between sessions. Your $HOME directory is private to you and cannot be accessed by other users.$HOME is not ~/bin. So 'C' is the answer.cd ~/ is egual at cd $HOME ~/bin is egual a cd $HOME/bin the persistent disk in cloud shell is for $HOMEWon't the persistent disk be auto-delete enabled by default, whereby the work maybe lost. Would that not be sufficient reason to consider Cloud storage instead. ThanksThe virtual machine instance that backs your Cloud Shell session is not permanently allocated to a Cloud Shell session and terminates if the session is inactive for an hour. After the instance is terminated, any modifications that you made to it outside your $HOME are lost.Cloud Shell provisions 5 GB of free persistent disk storage mounted as your $HOMEMaybe also to mention is that ~/bin is located in the $HOME directorycd ~/ is egual at cd $HOME ~/bin is egual a cd $HOME/bin the persistent disk in cloud shell is for $HOMEAgree. A is right

upvoted 77 times

akoti 5 years, 9 months ago

$HOME is not ~/bin. So 'C' is the answer.

upvoted 1 times

akoti 5 years, 9 months ago

$HOME is not ~/bin. So 'C' is the answer.cd ~/ is egual at cd $HOME ~/bin is egual a cd $HOME/bin the persistent disk in cloud shell is for $HOME

upvoted 1 times

zanfo 4 years, 11 months ago

cd ~/ is egual at cd $HOME ~/bin is egual a cd $HOME/bin the persistent disk in cloud shell is for $HOME

upvoted 10 times

Shabje 6 years, 3 months ago

Won't the persistent disk be auto-delete enabled by default, whereby the work maybe lost. Would that not be sufficient reason to consider Cloud storage instead. ThanksThe virtual machine instance that backs your Cloud Shell session is not permanently allocated to a Cloud Shell session and terminates if the session is inactive for an hour. After the instance is terminated, any modifications that you made to it outside your $HOME are lost.Cloud Shell provisions 5 GB of free persistent disk storage mounted as your $HOME

upvoted 2 times

Shabje 6 years, 3 months ago

Won't the persistent disk be auto-delete enabled by default, whereby the work maybe lost. Would that not be sufficient reason to consider Cloud storage instead. Thanks

upvoted 2 times

kaush 6 years, 1 month ago

The virtual machine instance that backs your Cloud Shell session is not permanently allocated to a Cloud Shell session and terminates if the session is inactive for an hour. After the instance is terminated, any modifications that you made to it outside your $HOME are lost.

upvoted 3 times

zanfo 4 years, 11 months ago

Cloud Shell provisions 5 GB of free persistent disk storage mounted as your $HOME

upvoted 1 times

Jambalaja 5 years, 4 months ago

Maybe also to mention is that ~/bin is located in the $HOME directorycd ~/ is egual at cd $HOME ~/bin is egual a cd $HOME/bin the persistent disk in cloud shell is for $HOMEAgree. A is right

upvoted 18 times

Jambalaja 5 years, 4 months ago

Maybe also to mention is that ~/bin is located in the $HOME directory

upvoted 18 times

AzureDP900 3 years, 10 months ago

Agree. A is right

upvoted 1 times

Eroc 6 years, 9 months ago

Well, I just double checked and if they were referring to the PATH variable then /usr/local/bin is also a correct answer.....

upvoted 21 times

izzetcanyc Most Recent 1 year, 4 months ago

The answer should be A because its standard location for user's scripts and custom utilities. D is also available for other sessions but /usr/local/bin is not suitable for custom utilities purposes also not all users have permission to this location. This location is for softwares for system wide.

upvoted 1 times

izekc 1 year, 5 months ago

~/bin will not work for all sessions. Only /usr/local/bin is more accurate

upvoted 1 times

jobolesonihal 1 year, 5 months ago

Warning: If you delete your $HOME directory, or if you don't access Cloud Shell for 120 days, Cloud Shell automatically deletes your $HOME directory and can't recover your files. You'll receive an email notification before this occurs. Starting a Cloud Shell session prevents the $HOME directory persistent storage from being recycled.

upvoted 1 times

RobertArnaud 1 year, 8 months ago

If "few weeks" means less than 120 days

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Cloud Shell's home directory: ~/bin is a subdirectory within your Cloud Shell home directory. Files placed in this directory are automatically added to your PATH environment variable, making them accessible from anywhere in your Cloud Shell session. 2. Persistence: Your Cloud Shell home directory is persistent across sessions. This means any files you store there, including those in ~/bin, will remain available even if you close and reopen Cloud Shell. 3. Convenience: This option is the most straightforward. It doesn't require any extra configuration or interaction with other services. Why D i snot correct? because /usr/local/bin: is a common directory for system-wide binaries, but in Cloud Shell, your home directory's

upvoted 5 times

thewalker 2 years, 9 months ago

/usr/local/bin is the place where the files will persist across sessions. Hence, D.

upvoted 2 times

heretolearnazure 2 years, 12 months ago

A is correct!

upvoted 1 times

moota 3 years, 6 months ago

I tested this. Although ~/bin is not in the default $PATH, choice D is definitely not persisting across sessions.1. When logging in to cloud shell for the first time, the ~/bin directory does not exist 2. mkdir ~/bin 3. After re-login to the cloud shell, $PATH will automatically add ~/bin So A is the correct answer

upvoted 2 times

moota 3 years, 6 months ago

I tested this. Although ~/bin is not in the default $PATH, choice D is definitely not persisting across sessions.

upvoted 2 times

Flight1976 3 years, 2 months ago

1. When logging in to cloud shell for the first time, the ~/bin directory does not exist 2. mkdir ~/bin 3. After re-login to the cloud shell, $PATH will automatically add ~/bin So A is the correct answer

upvoted 4 times

FI22 3 years, 7 months ago

At this moment default directory cant be set as Cloud storage bucket, so no C. A will be correct as zonal PD with preinstalled tools 5gb available that does not timeout!

upvoted 2 times

omermahgoub 3 years, 8 months ago

The recommended location for storing a custom utility file that you want to use in Cloud Shell and that should be in the default execution path and persist across sessions is option A: ~/bin.

The ~/bin directory is a personal directory that is in the default execution path for all users in Cloud Shell. Any executable files that you place in this directory will be available to you whenever you log in to Cloud Shell, and they will persist across sessions.

upvoted 4 times

omermahgoub 3 years, 8 months ago

The recommended location for storing a custom utility file that you want to use in Cloud Shell and that should be in the default execution path and persist across sessions is option A: ~/bin. The ~/bin directory is a personal directory that is in the default execution path for all users in Cloud Shell. Any executable files that you place in this directory will be available to you whenever you log in to Cloud Shell, and they will persist across sessions.Option B: Cloud Storage is not a suitable location for storing a custom utility file that you want to use in Cloud Shell, as it is not in the default execution path and would require additional steps to make it accessible. Option C: The /google/scripts directory is not a suitable location for storing a custom utility file, as it is not in the default execution path and is intended for use by Google Cloud system processes. Option D: The /usr/local/bin directory is a system directory that is in the default execution path for all users, but it is not a suitable location for storing a custom utility file, as any files that you place in this directory may be deleted or overwritten during system updates.

upvoted 4 times

omermahgoub 3 years, 8 months ago

Option B: Cloud Storage is not a suitable location for storing a custom utility file that you want to use in Cloud Shell, as it is not in the default execution path and would require additional steps to make it accessible. Option C: The /google/scripts directory is not a suitable location for storing a custom utility file, as it is not in the default execution path and is intended for use by Google Cloud system processes. Option D: The /usr/local/bin directory is a system directory that is in the default execution path for all users, but it is not a suitable location for storing a custom utility file, as any files that you place in this directory may be deleted or overwritten during system updates.

upvoted 3 times

Jailbreaker 3 years, 9 months ago

For sure correct answer is A

upvoted 1 times

megumin 3 years, 9 months ago

ok for A

upvoted 1 times

minmin2020 3 years, 10 months ago

A. ~/bin

upvoted 1 times

vpatiltech 4 years, 6 months ago

Cloud Shell provisions 5 GB of persistent disk storage mounted as your $HOME directory on the Cloud Shell instance. All files you store in your home directory, including scripts and user configuration files like .bashrc and .vimrc, persist between sessions. Reference- https://cloud.google.com/shell/?utm_source=google&utm_medium=cpc&utm_campaign=japac-IN-all-en-dr-bkwsrmkt-all-all-trial-e-dr-1009882&utm_content=text-ad-none-none-DEV_c-CRE_442449534611-ADGP_Hybrid%20%7C%20BKWS%20-%20EXA%20%7C%20Txt%20~%20Management%20Tools%20~%20Cloud%20Shell_cloud%20shell-general%20-%20Products-KWID_43700054972141701-kwd-837034669893&userloc_9302140-network_g&utm_term=KW_gcp%20cloud%20shell&gclsrc=ds&gclsrc=ds

upvoted 4 times

OrangeTiger 4 years, 7 months ago

I think D is correct.ummm

upvoted 2 times

Question #43 *Topic 1*

You want to create a private connection between your instances on Compute Engine and your on-premises data center. You require a connection of at least 20
Gbps. You want to follow Google-recommended practices. How should you set up the connection?

A. Create a VPC and connect it to your on-premises data center using Dedicated Interconnect.

B. Create a VPC and connect it to your on-premises data center using a single Cloud VPN.

C. Create a Cloud Content Delivery Network (Cloud CDN) and connect it to your on-premises data center using Dedicated Interconnect.

D. Create a Cloud Content Delivery Network (Cloud CDN) and connect it to your on-premises datacenter using a single Cloud VPN.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (100%)

AWS56 Highly Voted 6 years, 7 months ago
Cloud VPN supports unto 3 Gbps where as Interconnect can support unto 100 gbps... I'll go with AA is okhttps://cloud.google.com/network-connectivity/docs/vpn/concepts/overview#network-bandwidthA, 20Gbps dedicated interconnect is the way.A is required for consistent speed and VPN not supports that speed
upvoted 45 times

AWS56 Highly Voted 6 years, 7 months ago
Cloud VPN supports unto 3 Gbps where as Interconnect can support unto 100 gbps... I'll go with A
upvoted 45 times

tartar 6 years ago
A is ok
upvoted 7 times

fraloca 5 years, 7 months ago
https://cloud.google.com/network-connectivity/docs/vpn/concepts/overview#network-bandwidth
upvoted 3 times

nitinz 5 years, 5 months ago
A, 20Gbps dedicated interconnect is the way.
upvoted 2 times

AzureDP900 3 years, 10 months ago
A is required for consistent speed and VPN not supports that speed
upvoted 2 times

omermahgoub 3 years, 8 months ago
Answer is A: Dedicated Interconnect is a service that allows you to create a dedicated, high-bandwidth network connection between your on-premises data center and Google Cloud. It is the recommended solution for creating a private connection between your on-premises data center and Google Cloud when you require a connection of at least 20 Gbps.

Option B: Using a single Cloud VPN to connect your VPC to your on-premises data center is not suitable for a connection of at least 20 Gbps, as Cloud VPN has a maximum capacity of 30 Gbps.

Option C: The Cloud Content Delivery Network (Cloud CDN) is a globally distributed network of caching servers that speeds up the delivery of static and dynamic web content. It is not suitable for creating a private connection between your instances on Compute Engine and your on-premises data center.

Option D: Connecting your Cloud CDN to your on-premises data center using a single Cloud VPN is not suitable for a connection of at least 20 Gbps, as Cloud VPN has a maximum capacity of 30 Gbps.
upvoted 14 times

omermahgoub 3 years, 8 months ago

Answer is A: Dedicated Interconnect is a service that allows you to create a dedicated, high-bandwidth network connection between your on-premises data center and Google Cloud. It is the recommended solution for creating a private connection between your on-premises data center and Google Cloud when you require a connection of at least 20 Gbps. Option B: Using a single Cloud VPN to connect your VPC to your on-premises data center is not suitable for a connection of at least 20 Gbps, as Cloud VPN has a maximum capacity of 30 Gbps. Option C: The Cloud Content Delivery Network (Cloud CDN) is a globally distributed network of caching servers that speeds up the delivery of static and dynamic web content. It is not suitable for creating a private connection between your instances on Compute Engine and your on-premises data center. Option D: Connecting your Cloud CDN to your on-premises data center using a single Cloud VPN is not suitable for a connection of at least 20 Gbps, as Cloud VPN has a maximum capacity of 30 Gbps.Very nice answer, I think you meant 3 Gbps for Cloud VPN.

upvoted 14 times

MJCLOUD 3 years, 5 months ago

Very nice answer, I think you meant 3 Gbps for Cloud VPN.

upvoted 4 times

Ekramy_Elnaggar Most Recent 1 year, 9 months ago

1. Dedicated Interconnect for high bandwidth: Dedicated Interconnect is designed for high-bandwidth, private connections between your on-premises network and Google Virtual Private Cloud (VPC). It offers speeds of 10 Gbps up to 100 Gbps, meeting your requirement of at least 20 Gbps. 2. Private and secure connection: Dedicated Interconnect provides a direct physical connection, bypassing the public internet. This ensures a secure and private connection for your sensitive data. 3. Google-recommended practice: For demanding workloads that require high bandwidth and low latency, Google recommends Dedicated Interconnect over VPN.

upvoted 3 times

heretolearnazure 2 years, 12 months ago

high bandwidth means A

upvoted 1 times

greyhats13 3 years, 8 months ago

the question mention 20gbs for the least, it should be Dedicated Interconnect. The answer is A

upvoted 1 times

megumin 3 years, 9 months ago

ok for A

upvoted 1 times

Jay_Krish 3 years, 11 months ago

Any connection between On-Prem and GCP and requires high speed I'd choose dedicated interconnect

upvoted 1 times

avinashvidyarthi 4 years, 3 months ago

A is correct

upvoted 1 times

vincy2202 4 years, 7 months ago

A is the correct answer

upvoted 2 times

haroldbenites 4 years, 8 months ago

Go for A

upvoted 2 times

joe2211 4 years, 8 months ago

vote A

upvoted 2 times

duocnh 4 years, 8 months ago

vote A

upvoted 3 times

unnikrisb 4 years, 10 months ago

A is good option (easily elminate C & D) and B with connection speed. 10Gbps per link for Dedicated Interconnect and Direct Peering 1.5-3Gbps per tunnel for Cloud VPN 50Mbps to 10Gbps per connection - Partner Interconnect noSLA - Carrier Peering

upvoted 1 times

amxexam 4 years, 11 months ago

Let's go with option elimination A. Create a VPC and connect it to your on-premises data centre using Dedicated Interconnect. >> Is the only remaining best option after elimination. As per the document, its partner interconnects with VPN. Interconnect is between GCP and on-prem. URL1 B. Create a VPC and connect it to your on-premises data centre using a single Cloud VPN. >> max 3 gigabits per second (Gbps) eliminate the option. C. Create a Cloud Content Delivery Network (Cloud CDN) and connect it to your on-premises data centre using Dedicated Interconnect. >> CDN is for

egress traffic or static content hosting hence eliminate the option URL2 D. Create a Cloud Content Delivery Network (Cloud CDN) and connect it to your on-premises datacenter using a single Cloud VPN. >> CDN is for egress traffic or static content hosting hence eliminate the option URL2 Hence A URL1 https://cloud.google.com/network-connectivity/docs/interconnect/concepts/overview URL2 https://cloud.google.com/network-connectivity/docs/cdn-interconnect

upvoted 3 times

MamthaSJ 5 years, 1 month ago

Answer is A

upvoted 1 times

aviratna 5 years, 1 month ago

A is correct. Dedicated Interconnect supports upto 80 GBPS

upvoted 1 times

victory108 5 years, 3 months ago

A. Create a VPC and connect it to your on-premises data center using Dedicated Interconnect.

upvoted 2 times

Question #44 *Topic 1*

You are analyzing and defining business processes to support your startup's trial usage of GCP, and you don't yet know what consumer demand for your product will be. Your manager requires you to minimize GCP service costs and adhere to Google best practices. What should you do?

A. Utilize free tier and sustained use discounts. Provision a staff position for service cost management.

B. Utilize free tier and sustained use discounts. Provide training to the team about service cost management.

C. Utilize free tier and committed use discounts. Provision a staff position for service cost management.

D. Utilize free tier and committed use discounts. Provide training to the team about service cost management.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (100%)

ehgm Highly Voted 4 years, 7 months ago
Sustained are automatic discounts for running specific GCE a significant portion of the billing month: https://cloud.google.com/compute/docs/sustained-use-discounts Committed is for workloads with predictable resource needs between 1 year or 3 year, discount is up to 57% for most resources: https://cloud.google.com/compute/docs/instances/signing-up-committed-use-discountsBest answer!Adding to it. Provide training to the team about service cost management: Because we are still using free tier, and no need a separate position even if that need arises. Its startup Company any way, and they are trained to work beyond hours :-)
upvoted 45 times

ehgm Highly Voted 4 years, 7 months ago
Sustained are automatic discounts for running specific GCE a significant portion of the billing month: https://cloud.google.com/compute/docs/sustained-use-discounts

Committed is for workloads with predictable resource needs between 1 year or 3 year, discount is up to 57% for most resources: https://cloud.google.com/compute/docs/instances/signing-up-committed-use-discounts
upvoted 45 times

squishy_fishy 4 years, 6 months ago
Best answer!
upvoted 5 times

ashishdwi007 2 years, 6 months ago
Adding to it. Provide training to the team about service cost management: Because we are still using free tier, and no need a separate position even if that need arises. Its startup Company any way, and they are trained to work beyond hours :-)
upvoted 2 times

crypt0 6 years, 9 months ago
I would choose "B"
upvoted 40 times

crypt0 6 years, 9 months ago
I would choose "B"B is okB reason minimize GCP service costs
upvoted 40 times

tartar 6 years ago
B is ok
upvoted 11 times

nitinz 5 years, 5 months ago
B reason minimize GCP service costs
upvoted 3 times

Ekramy_Elnaggar Most Recent 1 year, 9 months ago
1. Free Tier: Google Cloud's Free Tier allows you to use certain services for free, up to specified limits. This is ideal for a startup testing the platform and controlling costs during the trial phase. 2. Sustained Use Discounts: These discounts are automatically applied to your bill when you use eligible services for a significant portion of the billing month. This is beneficial for a startup that might have unpredictable usage patterns during the trial

phase, as it doesn't require any upfront commitment. 3. Training for cost management: Educating your team about cost management best practices empowers them to make cost-effective decisions from the start. This includes understanding pricing models, monitoring resource usage, and optimizing configurations.

upvoted 4 times

omermahgoub 3 years, 8 months ago

The recommended approach for minimizing GCP service costs and adhering to Google best practices are: - Free tier: Google Cloud offers a free tier of services that allows you to try out many of its products for free, up to certain usage limits. Utilizing the free tier can help you minimize your GCP service costs while you are in the trial usage phase. - Committed use discounts: Committed use discounts are a type of discount that you can apply to certain GCP products by committing to a certain level of usage over a one- or three-year period. Committed use discounts can help you save on your GCP service costs, but they require you to commit to a certain level of usage, which may not be suitable if you are unsure of your future demand. - Providing training to the team about service cost management: It is important that your team is aware of the different options available for minimizing GCP service costs and understands how to manage and monitor their usage of GCP services. Providing training on service cost management can help your team make informed decisions about how to use GCP services in the most cost-effective way.The recommended approach for minimizing GCP service costs and adhering to Google best practices while your startup is in the trial usage phase and you don't yet know what consumer demand for your product will be is option D: Utilize free tier and committed use discounts. Provide training to the team about service cost management.Sustained use discounts are based on your usage of GCP services over a certain period, and are not available for all GCP products. Provisioning a staff position for service cost management may not be necessary if you provide training to the team about service cost management.

upvoted 9 times

omermahgoub 3 years, 8 months ago

The recommended approach for minimizing GCP service costs and adhering to Google best practices are: - Free tier: Google Cloud offers a free tier of services that allows you to try out many of its products for free, up to certain usage limits. Utilizing the free tier can help you minimize your GCP service costs while you are in the trial usage phase.

- Committed use discounts: Committed use discounts are a type of discount that you can apply to certain GCP products by committing to a certain level of usage over a one- or three-year period. Committed use discounts can help you save on your GCP service costs, but they require you to commit to a certain level of usage, which may not be suitable if you are unsure of your future demand.

- Providing training to the team about service cost management: It is important that your team is aware of the different options available for minimizing GCP service costs and understands how to manage and monitor their usage of GCP services. Providing training on service cost management can help your team make informed decisions about how to use GCP services in the most cost-effective way.

upvoted 9 times

omermahgoub 3 years, 8 months ago

The recommended approach for minimizing GCP service costs and adhering to Google best practices while your startup is in the trial usage phase and you don't yet know what consumer demand for your product will be is option D: Utilize free tier and committed use discounts. Provide training to the team about service cost management.

upvoted 1 times

omermahgoub 3 years, 8 months ago

Sustained use discounts are based on your usage of GCP services over a certain period, and are not available for all GCP products. Provisioning a staff position for service cost management may not be necessary if you provide training to the team about service cost management.

upvoted 2 times

SureshbabuK 3 years, 8 months ago

Sustained use discounts - Compute Engine - Google Cloudhttps://cloud.google.com › ... › Documentation Compute Engine offers sustained use discounts on resources that are used for more than 25% of a billing month - There trial could be more than 7 or 8 days , at this point commitment of use can not be provided due to trails stage of gcp use

upvoted 1 times

vranjan 3 years, 9 months ago

The answer is B, because Sustained use discount can give up to 30% and requires no commitment.

upvoted 1 times

megumin 3 years, 9 months ago

ok for B

upvoted 1 times

zr79 3 years, 10 months ago

committed use discounts are for long-run discounts which in the case of startup they're trying GCP. So options C and D are out B is the correct answer

upvoted 1 times

AzureDP900 3 years, 10 months ago

I will choose B, D is only long term commitment

upvoted 1 times

minmin2020 3 years, 10 months ago

B. Utilize free tier and sustained use discounts. Provide training to the team about service cost management.

upvoted 2 times

BiddlyBdoyng 3 years, 10 months ago

Sustained use discount makes sense over committed as not enough info to know what to comit to. Provide staff training on how to keep things cheap gonna further keep cost down.

upvoted 1 times

cmamiusa 4 years, 4 months ago

B is the correct option

upvoted 1 times

gcmrjbr 4 years, 5 months ago

free tier (monthly discounts) does not make sense combined with committed use discounts - anual base, dont't you think so?

upvoted 2 times

vincy2202 4 years, 7 months ago

B is the correct answer

upvoted 1 times

Israel 4 years, 11 months ago

Answer is B

upvoted 1 times

VishalB 5 years ago

Answer B Sustained use discounts are applied on incremental use after you reach certain usage thresholds. This means that you pay only for the number of minutes that you use an instance, and Compute Engine automatically gives you the best price. There's no reason to run an instance for longer than you need it. - https://cloud.google.com/compute/docs/sustained-use-discounts Committed use discounts are ideal for workloads with predictable resource needs. When you purchase a committed use contract, you purchase compute resource (vCPUs, memory, GPUs, and local SSDs) at a discounted price in return for committing to paying for those resources for 1 year or 3 years. The discount is up to 57% for most resources like machine types or GPUs. The discount is up to 70% for memory-optimized machine types. For committed use prices for different machine types, see VM instances pricing. - https://cloud.google.com/compute/docs/instances/signing-up-committed-use-discounts

upvoted 5 times

MamthaSJ 5 years, 1 month ago

Answer is B

upvoted 1 times

Question #45 *Topic 1*

You are building a continuous deployment pipeline for a project stored in a Git source repository and want to ensure that code changes can be verified before deploying to production. What should you do?

A. Use Spinnaker to deploy builds to production using the red/black deployment strategy so that changes can easily be rolled back.

B. Use Spinnaker to deploy builds to production and run tests on production deployments.

C. Use Jenkins to build the staging branches and the master branch. Build and deploy changes to production for 10% of users before doing a complete rollout.

D. Use Jenkins to monitor tags in the repository. Deploy staging tags to a staging environment for testing. After testing, tag the repository for production and deploy that to the production environment.

Suggested Answer: *D*

Reference: https://github.com/GoogleCloudPlatform/continuous-deployment-on-kubernetes/blob/master/README.md

Community vote distribution

D (96%) (4

Googler2 Highly Voted 6 years, 4 months ago

I believe the best answer is D, because the tagging is a best practice that is recommended on Jenkins/Spinnaker to deploy the right code and prevent accidentally (or intentionally) push of wrong code to production environments. See https://stackify.com/continuous-delivery-git-jenkins/Agreed with D as the right answer. The url provided explains it betteryes, D is correctHow can I join Google

upvoted 61 times

Googler2 Highly Voted 6 years, 4 months ago

I believe the best answer is D, because the tagging is a best practice that is recommended on Jenkins/Spinnaker to deploy the right code and prevent accidentally (or intentionally) push of wrong code to production environments. See https://stackify.com/continuous-delivery-git-jenkins/

upvoted 61 times

Ziegler 6 years, 2 months ago

Agreed with D as the right answer. The url provided explains it better

upvoted 10 times

AzureDP900 3 years, 10 months ago

yes, D is correct

upvoted 1 times

zr79 3 years, 10 months ago

How can I join Google

upvoted 6 times

Anish17 5 years, 12 months ago

I got this question in real exam. This question states "before deploying to production" environment. So i picked D . I passed the exam.that resolved the puzzle :)congrats for passing the exam. practising all 255 questions is sufficient for passing the exam? how much percentage of questions you got from here roughly?He won't answer, he already passed the examAgree, its the only answer that meets the requirement of "before deploying to production"not only 1 Q passed! C is beeter

upvoted 53 times

Anish17 5 years, 12 months ago

I got this question in real exam. This question states "before deploying to production" environment. So i picked D . I passed the exam.

upvoted 53 times

bnlcnd 5 years, 6 months ago

that resolved the puzzle :)

upvoted 6 times

GunaGCP123 4 years, 9 months ago

congrats for passing the exam. practising all 255 questions is sufficient for passing the exam? how much percentage of questions you got from here roughly?

upvoted 2 times

GunaGCP123 4 years, 9 months ago

congrats for passing the exam. practising all 255 questions is sufficient for passing the exam? how much percentage of questions you got from here roughly?He won't answer, he already passed the exam

upvoted 2 times

zr79 3 years, 10 months ago

He won't answer, he already passed the exam

upvoted 10 times

Rzla 4 years, 11 months ago

Agree, its the only answer that meets the requirement of "before deploying to production"

upvoted 2 times

winset 5 years, 6 months ago

not only 1 Q passed! C is beeter

upvoted 2 times

khadar2001 Most Recent 1 year, 4 months ago

C - can be one of the option. Verified from practice test

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. Clear separation of environments: Using separate staging and production environments ensures that code changes are thoroughly tested before they reach your users. This minimizes the risk of introducing bugs or regressions into production. 2. Controlled deployments: Tagging the repository for different stages (staging, production) provides a clear and auditable way to track which code versions are deployed in each environment. This makes it easier to roll back to a previous version if necessary. 3. Jenkins for automation: Jenkins is a powerful automation server that can monitor your Git repository for new tags, automatically build and deploy the code to the appropriate environment, and even trigger automated tests. 4. Comprehensive testing: A staging environment allows you to perform comprehensive testing, including integration testing, user acceptance testing (UAT), and performance testing, before deploying to production.

upvoted 3 times

sim7243 1 year, 9 months ago

D is the best choice.

upvoted 1 times

bkovari 2 years, 5 months ago

ABC is about to deploy to production without prior testing. Hence D is the only reasonable choice.

upvoted 1 times

ashishdwi007 2 years, 6 months ago

A B and C are very risky options and irrelevant as involving production. Whereas question is to test before rolling out to production. Hence D

upvoted 1 times

blackhawk86 2 years, 11 months ago

No discussion here.

upvoted 1 times

megumin 3 years, 9 months ago

ok for D

upvoted 1 times

BiddlyBdoyng 3 years, 10 months ago

Given the question D is the only answer. Everything else pushes to production immediately.

upvoted 1 times

Davidik79 4 years, 5 months ago

The question states: "... code changes can be verified BEFORE deploying to production", it eliminates option C. The approach of tagging is the correct practise that DevOps use

upvoted 2 times

hantanbl 4 years, 7 months ago

C is the closest answer If question is asking 'what's Jenkin best practise' then D is the answerC involves deploying into production. the question specifies: "BEFORE deploying to production"Incorrect. D is the correct answer. 1. It should test BEFORE deploy to production. 2. It's a DevOps practice

upvoted 1 times

hantanbl 4 years, 7 months ago

C is the closest answer If question is asking 'what's Jenkin best practise' then D is the answer

upvoted 1 times

Davidik79 4 years, 5 months ago

C involves deploying into production. the question specifies: "BEFORE deploying to production"

upvoted 2 times

blackhawk86 2 years, 11 months ago

Incorrect. D is the correct answer. 1. It should test BEFORE deploy to production. 2. It's a DevOps practice

upvoted 1 times

lxgywil 4 years, 7 months ago

I choose D as we want to ensure that code changes can be verified BEFORE deploying to production. Option C suggests that we build and deploy changes to production for 10% of users.

upvoted 1 times

OrangeTiger 4 years, 7 months ago

Vote D. C is canary deploy. But the sentence in question has no word to mean "tested by a small number of users"The revelal solution on this site is wrong, isn't it? I'm getting anxious.

upvoted 3 times

OrangeTiger 4 years, 7 months ago

Selected Answer: D

Vote D. C is canary deploy. But the sentence in question has no word to mean "tested by a small number of users"

upvoted 3 times

OrangeTiger 4 years, 7 months ago

The revelal solution on this site is wrong, isn't it? I'm getting anxious.

upvoted 1 times

vincy2202 4 years, 7 months ago

D is the correct answer

upvoted 1 times

ABO_Doma 4 years, 8 months ago

clearly

upvoted 2 times

phantomsg 4 years, 8 months ago

Question asks to verify before Production. So tagging and deploying appropriately is the right approach.

upvoted 3 times

Question #46 *Topic 1*

You have an outage in your Compute Engine managed instance group: all instances keep restarting after 5 seconds. You have a health check configured, but autoscaling is disabled. Your colleague, who is a Linux expert, offered to look into the issue. You need to make sure that he can access the VMs. What should you do?

A. Grant your colleague the IAM role of project Viewer

B. Perform a rolling restart on the instance group

C. Disable the health check for the instance group. Add his SSH key to the project-wide SSH Keys

D. Disable autoscaling for the instance group. Add his SSH key to the project-wide SSH Keys

Suggested Answer: *C*

Official Answer: C

Community vote distribution

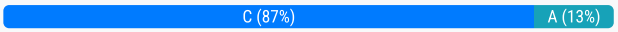

Narinder Highly Voted 4 years, 7 months ago

C, is the correct answer. As per the requirement linux expert would need access to VM to troubleshoot the issue. With health check enabled, old VM will be terminated as soon as health-check fails for the VM and new VM will be auto-created. So, this situation will prevent linux expert to troubleshoot the issue. Had it been the case that stack-drover logging is enabled and the expert just want to view the logs from the Cloud-logs than role to project-viewer could help. But it is specifically mentioned that expert will login into VM to troubleshoot the issue and not looking at the cloud Logs. So, Option-C is the correct answer.great answerThis looks best justification between A and C.. So, C should be correct answer.Good explanationThe key element in C is "Disable the Health check.", so that server wont restart automatically. But before that the actual troubleshooting step is to check Cloud console -> Instance template -> Metadata-> and see if any startup script is there, if yes review it and possibly remove it. [Consider the case, a script is causing restarting the VM, (possibly in Metadata).]

upvoted 99 times

Narinder Highly Voted 4 years, 7 months ago

C, is the correct answer. As per the requirement linux expert would need access to VM to troubleshoot the issue. With health check enabled, old VM will be terminated as soon as health-check fails for the VM and new VM will be auto-created. So, this situation will prevent linux expert to troubleshoot the issue. Had it been the case that stack-drover logging is enabled and the expert just want to view the logs from the Cloud-logs than role to project-viewer could help. But it is specifically mentioned that expert will login into VM to troubleshoot the issue and not looking at the cloud Logs. So, Option-C is the correct answer.

upvoted 99 times

twistyfries 4 years, 5 months ago

great answer

upvoted 2 times

AmitAr 4 years, 3 months ago

This looks best justification between A and C.. So, C should be correct answer.

upvoted 2 times

devjuliusoh 3 years, 12 months ago

Good explanation

upvoted 2 times

ashishdwi007 2 years, 6 months ago

The key element in C is "Disable the Health check.", so that server wont restart automatically. But before that the actual troubleshooting step is to check Cloud console -> Instance template -> Metadata-> and see if any startup script is there, if yes review it and possibly remove it. [Consider the case, a script is causing restarting the VM, (possibly in Metadata).]

upvoted 1 times

KouShikyou 6 years, 10 months ago

C should be correct answer.

upvoted 39 times

Ekramy_Elnaggar Most Recent 1 year, 9 months ago

1. Access for troubleshooting: Disabling the health check prevents the VMs from constantly restarting, allowing the expert to log in and investigate the root cause. 2. SSH access: Adding the expert's SSH key grants them the necessary access to the VMs. 3. Temporary measure: This is a temporary measure for troubleshooting. Once the issue is resolved, the health check should be re-enabled.

upvoted 3 times

wilson1005 2 years, 5 months ago

C is correct!

upvoted 1 times

heretolearnazure 2 years, 12 months ago

I like C

upvoted 1 times

JC0926 3 years, 5 months ago

To allow your colleague, who is a Linux expert, to access the VMs and troubleshoot the issue, you should disable the health check for the instance group. This will prevent the instance group from automatically removing and replacing unhealthy instances. You should also add your colleague's SSH key to the project-wide SSH Keys to allow him to SSH into the instances and perform troubleshooting. This can be done through the Google Cloud Console or the gcloud command-line tool.

upvoted 4 times

MestreCholas 3 years, 5 months ago

C. Disable the health check for the instance group. Add his SSH key to the project-wide SSH Keys. Granting the IAM role of project Viewer (Option A) would allow your colleague to view the project resources but not necessarily give them access to the specific VM instances. Performing a rolling restart on the instance group (Option B) would not resolve the issue, as the instances keep restarting after 5 seconds. Disabling autoscaling for the instance group (Option D) is not relevant since autoscaling is already disabled. Disabling the health check for the instance group will prevent the managed instance group from automatically recreating the instances. Adding your colleague's SSH key to the project-wide SSH Keys will allow them to access the VMs and troubleshoot the issue. Once the issue is resolved, you can re-enable the health check for the instance group.

upvoted 6 times

urssiva 3 years, 6 months ago

C is the answer

upvoted 1 times

roaming_panda 3 years, 7 months ago

C as the machines will keep restarting if hc is not disabled , and then our expert can look around for RCA

upvoted 1 times

omermahgoub 3 years, 8 months ago

To allow your colleague access the instances in MIG you should do option D: Disable autoscaling for the instance group. Add his SSH key to the project-wide SSH Keys.

Disabling autoscaling for the instance group will prevent new instances from being created or terminated while your colleague is trying to troubleshoot the issue. This will give him a stable environment to work in and will ensure that he can access the instances that are currently in the instance group.

Adding his SSH key to the project-wide SSH Keys will allow him to connect to the instances using Secure Shell (SSH) without requiring a password. This is a convenient way to give him access to the instances and can help him troubleshoot the issue more efficiently.

upvoted 2 times

omermahgoub 3 years, 8 months ago

To allow your colleague access the instances in MIG you should do option D: Disable autoscaling for the instance group. Add his SSH key to the project-wide SSH Keys. Disabling autoscaling for the instance group will prevent new instances from being created or terminated while your colleague is trying to troubleshoot the issue. This will give him a stable environment to work in and will ensure that he can access the instances that are currently in the instance group. Adding his SSH key to the project-wide SSH Keys will allow him to connect to the instances using Secure Shell (SSH) without requiring a password. This is a convenient way to give him access to the instances and can help him troubleshoot the issue more efficiently.autoscaling is already disabled. Answer D is not make sense here. i voted it wrongly instead of reply.Option A: Granting your colleague the IAM role of project Viewer will not allow him to access the instances in the managed instance group. The project Viewer role does not include any permissions to access Compute Engine resources. Option B: Performing a rolling restart on the instance group will not solve the issue and may even make it worse if the instances are not able to start up properly. Option C: Disabling the health check for the instance group will not solve the issue and may even make it worse if the instances are not able to start up properly. Adding his SSH key to the project-wide SSH Keys will allow him to access the instances, but it is not a sufficient solution on its own.

upvoted 2 times

n_nana 3 years, 7 months ago

autoscaling is already disabled. Answer D is not make sense here. i voted it wrongly instead of reply.

upvoted 1 times

omermahgoub 3 years, 8 months ago

Option A: Granting your colleague the IAM role of project Viewer will not allow him to access the instances in the managed instance group. The project Viewer role does not include any permissions to access Compute Engine resources. Option B: Performing a rolling restart on the instance group will not solve the issue and may even make it worse if the instances are not able to start up properly. Option C: Disabling the health check for the instance group will not solve the issue and may even make it worse if the instances are not able to start up properly. Adding his SSH key to the project-wide SSH Keys will allow him to access the instances, but it is not a sufficient solution on its own.

upvoted 2 times

surajkrishnamurthy 3 years, 8 months ago

C is the correct answer

upvoted 1 times

megumin 3 years, 9 months ago

ok for C

upvoted 1 times

minmin2020 3 years, 10 months ago

C. Disable the health check for the instance group. Add his SSH key to the project-wide SSH Keys. This will allow the engineer to logon to the VM's and check the logs. Disabling health checks will prevent rebooting of the VM's.

upvoted 2 times

holerina 3 years, 11 months ago

C should be the correct answer

upvoted 1 times

abirroy 3 years, 11 months ago

Disable the health check for the instance group. Add his SSH key to the project-wide SSH Keys

upvoted 1 times

backhand 4 years ago

vote C what can project viewer do without access vm by linux expert?

upvoted 1 times

tricky_learner 4 years, 1 month ago

I believe that answer C is correct.

upvoted 1 times

Question #47 *Topic 1*

Your company is migrating its on-premises data center into the cloud. As part of the migration, you want to integrate Google Kubernetes Engine (GKE) for workload orchestration. Parts of your architecture must also be PCI DSS-compliant. Which of the following is most accurate?

A. App Engine is the only compute platform on GCP that is certified for PCI DSS hosting.

B. GKE cannot be used under PCI DSS because it is considered shared hosting.

C. GKE and GCP provide the tools you need to build a PCI DSS-compliant environment.

D. All Google Cloud services are usable because Google Cloud Platform is certified PCI-compliant.

Suggested Answer: *C*

Official Answer: C

Community vote distribution

C (100%)

rishab86 Highly Voted 5 years, 2 months ago

Link : https://cloud.google.com/security/compliance/pci-dss Clearly mention GKE as PCI DSS-Compliant but not all GCP service are PCI DSS-Compliant so answer is definitely C.C – Kubernetes Engine provides tools you need to build to PCI-DSS compliant environment.But, The paragraph 3 says that all products of google are certified by PCI.In 2022, GCP is now fully PCI-DSS compliant, so technically D is perfectly true. But you still have to check that your application is PCI-DSS compliant. so C is still the best answer.

upvoted 48 times

rishab86 Highly Voted 5 years, 2 months ago

Link : https://cloud.google.com/security/compliance/pci-dss Clearly mention GKE as PCI DSS-Compliant but not all GCP service are PCI DSS-Compliant so answer is definitely C.

upvoted 48 times

MaxNRG 4 years, 9 months ago

C – Kubernetes Engine provides tools you need to build to PCI-DSS compliant environment.

upvoted 1 times

haroldbenites 4 years, 8 months ago

But, The paragraph 3 says that all products of google are certified by PCI.

upvoted 2 times

Mikado211 4 years ago

In 2022, GCP is now fully PCI-DSS compliant, so technically D is perfectly true. But you still have to check that your application is PCI-DSS compliant. so C is still the best answer.

upvoted 8 times

aviratna 5 years, 1 month ago

C: GKE & Compute Engine is PCI DSS compliant while Cloud Function, App Engine are not PC compliant

upvoted 5 times

Ekramy_Elnaggar Most Recent 1 year, 9 months ago

1. GKE and PCI DSS: While GKE itself isn't inherently PCI DSS compliant, it provides the infrastructure and tools you need to build a compliant environment. You'll need to configure it correctly, implement security measures, and follow best practices. 2. Shared Responsibility Model: Google Cloud Platform operates under a shared responsibility model. Google is responsible for securing the underlying infrastructure, while you are responsible for securing your applications and data within that environment. 3. Flexibility for Compliance: GKE offers features like private clusters, network policies, and integration with security tools that help you meet PCI DSS requirements.

upvoted 3 times

eka_nostra 3 years ago

We still have to configure our env to comply with PCI/DSS. https://cloud.google.com/architecture/pci-dss-compliance-in-gcp#kubernetes_engine

upvoted 1 times

omermahgoub 3 years, 8 months ago

The most accurate statement is option C: GKE and GCP provide the tools you need to build a PCI DSS-compliant environment. Google Kubernetes Engine (GKE) is a fully managed service that allows you to deploy and manage containerized applications on Google Cloud. It is not specifically certified for PCI DSS hosting, but it can be used as part of a PCI DSS-compliant environment if the necessary controls and safeguards are in place.

Google Cloud Platform (GCP) provides a range of tools and services that can be used to build a PCI DSS-compliant environment, including Cloud Identity and Access Management (IAM) for controlling access to resources, Cloud Key Management Service (KMS) for managing encryption keys, and Cloud Security Command Center for monitoring and detecting security threats.Option A: App Engine is a fully managed platform for building and deploying web and mobile applications, but it is not the only compute platform on GCP that is certified for PCI DSS hosting. Other compute platforms such as Compute Engine and Google Kubernetes Engine can also be used as part of a PCI DSS-compliant environment. Option B: GKE is not considered shared hosting and can be used as part of a PCI DSS-compliant environment if the necessary controls and safeguards are in place. Option D: While Google Cloud Platform is certified PCI-compliant, not all of its services are automatically usable in a PCI DSS-compliant environment. It is up to the user to ensure that they are using the appropriate controls and safeguards to meet the requirements of the PCI DSS.

upvoted 4 times

omermahgoub 3 years, 8 months ago

The most accurate statement is option C: GKE and GCP provide the tools you need to build a PCI DSS-compliant environment.

Google Kubernetes Engine (GKE) is a fully managed service that allows you to deploy and manage containerized applications on Google Cloud. It is not specifically certified for PCI DSS hosting, but it can be used as part of a PCI DSS-compliant environment if the necessary controls and safeguards are in place.

Google Cloud Platform (GCP) provides a range of tools and services that can be used to build a PCI DSS-compliant environment, including Cloud Identity and Access Management (IAM) for controlling access to resources, Cloud Key Management Service (KMS) for managing encryption keys, and Cloud Security Command Center for monitoring and detecting security threats.

upvoted 4 times

omermahgoub 3 years, 8 months ago

Option A: App Engine is a fully managed platform for building and deploying web and mobile applications, but it is not the only compute platform on GCP that is certified for PCI DSS hosting. Other compute platforms such as Compute Engine and Google Kubernetes Engine can also be used as part of a PCI DSS-compliant environment. Option B: GKE is not considered shared hosting and can be used as part of a PCI DSS-compliant environment if the necessary controls and safeguards are in place. Option D: While Google Cloud Platform is certified PCI-compliant, not all of its services are automatically usable in a PCI DSS-compliant environment. It is up to the user to ensure that they are using the appropriate controls and safeguards to meet the requirements of the PCI DSS.

upvoted 3 times

abirroy 3 years, 11 months ago

C is the right answer

upvoted 1 times

vincy2202 4 years, 7 months ago

C is the correct answer

upvoted 1 times

haroldbenites 4 years, 8 months ago

Go for C.

upvoted 1 times

SHOURYA_SOOD 4 years, 8 months ago

C- All of them: GKE, GCE, and GAE ate PCI-DSS-Compliant but A & B says it's only GAE and GCE respectively so cancel them out. D says all of GCP is PCI DSS-Compliant but it's not true. So, C seems to be the right answer.

upvoted 1 times

imranmani 4 years, 10 months ago

C is the right answer

upvoted 1 times

MamthaSJ 5 years, 1 month ago

Answer is C

upvoted 3 times

victory108 5 years, 1 month ago

C. GKE and GCP provide the tools you need to build a PCI DSS-compliant environment.

upvoted 1 times

Question #48 *Topic 1*

Your company has multiple on-premises systems that serve as sources for reporting. The data has not been maintained well and has become degraded over time.
You want to use Google-recommended practices to detect anomalies in your company data. What should you do?

A. Upload your files into Cloud Storage. Use Cloud Datalab to explore and clean your data.

B. Upload your files into Cloud Storage. Use Cloud Dataprep to explore and clean your data.

C. Connect Cloud Datalab to your on-premises systems. Use Cloud Datalab to explore and clean your data.

D. Connect Cloud Dataprep to your on-premises systems. Use Cloud Dataprep to explore and clean your data.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (100%)

JohnWick2020 Highly Voted 5 years, 4 months ago
Answer is B: Keynotes from question: 1- On-premise data sources 2- Unfit data; not well maintained and degraded 3- Google-recommended best practice to "detect anomalies" <<-Very important. Explanation: A & C - incorrect; Datalab does not provide anomaly detection OOTB. It is used more for data science scenarios like interactive data analysis and build ML models. B - CORRECT; DataPrep OOTB provides for fast exploration and anomaly detection and lists cloud storage as an ingestion medium. Refer to ELT pipeline architecture here = https://cloud.google.com/dataprep D - incorrect; At this time DataPrep cannot connect to SaaS or on-premise source. Not to be confused for DataFlow which can!
upvoted 65 times

Eroc Highly Voted 6 years, 9 months ago
Both B and D work, because the question says "Google's Best Practices" uploading the files first would keep the original copies Google encrypted and stored.Both of them works....You can't connect DataPrep to your on-prem systems. You simply upload a file, but that is not connecting it to your systems. Because of that, I'd discard D and stay with B.B is okB, dataprep = visually explore, clean, and prepare data for analysisB is better choice
upvoted 12 times

Eroc 6 years, 9 months ago
Both B and D work, because the question says "Google's Best Practices" uploading the files first would keep the original copies Google encrypted and stored.
upvoted 12 times

skywalker 6 years, 3 months ago
Both of them works....You can't connect DataPrep to your on-prem systems. You simply upload a file, but that is not connecting it to your systems. Because of that, I'd discard D and stay with B.
upvoted 1 times

skywalker 6 years, 3 months ago
Both of them works....
upvoted 1 times

Musk 6 years ago
You can't connect DataPrep to your on-prem systems. You simply upload a file, but that is not connecting it to your systems. Because of that, I'd discard D and stay with B.
upvoted 9 times

tartar 6 years ago
B is ok
upvoted 9 times

nitinz 5 years, 5 months ago
B, dataprep = visually explore, clean, and prepare data for analysis
upvoted 7 times

AzureDP900 3 years, 10 months ago
B is better choice
upvoted 1 times

Johnb9682 Most Recent 4 months, 1 week ago

Althought Dataprep is deprecated

upvoted 1 times

Sib09 1 year, 5 months ago

you have to bring data to Cloud storage before running data prep.

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

Why Cloud Storage is important ? 1. Centralized repository: Cloud Storage provides a secure and scalable place to store your data. This makes it accessible to various GCP services. 2. Data lake concept: This aligns with the idea of a data lake, where you bring raw data into a central location before processing and refining it. Why Cloud Dataprep is a good fit ? 1. Visual data exploration: Dataprep excels at helping you quickly understand your data through visualizations and profiling. This is crucial for identifying anomalies. 2. Data cleaning and transformation: Dataprep makes it easy to clean and standardize your data, which is essential before anomaly detection. Inconsistent formats, missing values, and errors can skew your analysis. 3. Built-in anomaly detection: Dataprep has features specifically designed to help you find anomalies. It can highlight unusual values, outliers, and patterns.

upvoted 4 times

snehaso 2 years ago

Datalab was shutdown. Its replacement is vertex AI. Read question accordingly

upvoted 1 times

thewalker 2 years, 9 months ago

Cloud Datalab is a powerful interactive tool created to explore, analyze, transform, and visualize data and build machine learning models on Google Cloud Platform. Dataprep by Trifacta is an intelligent data service for visually exploring, cleaning, and preparing structured and unstructured data for analysis, reporting, and machine learning. Dataprep do not have an integration for on-prem: https://console.cloud.google.com/marketplace/product/endpoints/cloud-dataprep-editions-v2?project=fast-art-401415 So, clearly, the only option left is B.

upvoted 4 times

heretolearnazure 2 years, 12 months ago

B is correct.

upvoted 1 times

n_nana 3 years, 5 months ago

Today, data ingestion to DataPrep can be Application, file upload, database. so B is also now valid

upvoted 1 times

omermahgoub 3 years, 8 months ago

The recommended approach for detecting anomalies in your company data using Google-recommended practices is option B: Upload your files into Cloud Storage. Use Cloud Dataprep to explore and clean your data.

Cloud Storage is a highly scalable, durable, and secure object storage service that can be used to store and retrieve data from anywhere on the web. You can use Cloud Storage to store your company data files and make them available for analysis.

Cloud Dataprep is a fully managed data preparation service that allows you to quickly and easily explore, clean, and transform your data for analysis. It can help you detect anomalies in your data by providing features such as data profiling, data cleansing, and data transformation.

upvoted 2 times

omermahgoub 3 years, 8 months ago

The recommended approach for detecting anomalies in your company data using Google-recommended practices is option B: Upload your files into Cloud Storage. Use Cloud Dataprep to explore and clean your data. Cloud Storage is a highly scalable, durable, and secure object storage service that can be used to store and retrieve data from anywhere on the web. You can use Cloud Storage to store your company data files and make them available for analysis. Cloud Dataprep is a fully managed data preparation service that allows you to quickly and easily explore, clean, and transform your data for analysis. It can help you detect anomalies in your data by providing features such as data profiling, data cleansing, and data transformation.Option A: Using Cloud Datalab to explore and clean your data is not a recommended approach, as Cloud Datalab is a collaborative data exploration and visualization platform that is not specifically designed for data preparation tasks such as cleansing and transformation. Option C: Connecting Cloud Datalab to your on-premises systems is not a recommended approach, as Cloud Datalab is a collaborative data exploration and visualization platform and is not designed for data preparation tasks such as cleansing and transformation. Option D: Connecting Cloud Dataprep to your on-premises systems is not necessary, as you can use Cloud Dataprep to explore and clean data stored in Cloud Storage.

upvoted 2 times

omermahgoub 3 years, 8 months ago

Option A: Using Cloud Datalab to explore and clean your data is not a recommended approach, as Cloud Datalab is a collaborative data exploration and visualization platform that is not specifically designed for data preparation tasks such as cleansing and transformation. Option C: Connecting Cloud Datalab to your on-premises systems is not a recommended approach, as Cloud Datalab is a collaborative data exploration and visualization platform and is not designed for data preparation tasks such as cleansing and transformation. Option D: Connecting Cloud Dataprep to your on-premises systems is not necessary, as you can use Cloud Dataprep to explore and clean data stored in Cloud Storage.

upvoted 1 times

allen_y_q_huang 3 years, 8 months ago

ok for B & D, but B is suitable to gcp

upvoted 1 times

Smaks 3 years, 8 months ago

Datalab is deprecated : https://cloud.google.com/datalab/docs New Cloud Dataprep options will give connectivity to relational databases, business applications and extend our integrations across Google Cloud with Google Sheets: https://www.trifacta.com/blog/cloud-dataprep-trifacta/

upvoted 1 times

megumin 3 years, 9 months ago

ok for B

upvoted 1 times

Cloudexplorer 4 years ago

Could anyone provide a link where it explicitly says that Datprep does not connect to on-premises data sources. In the ingestion layer on the diagram at https://cloud.google.com/dataprep it shows databases as a source. I can't see anywhere that there is a limitation connecting to on-premises. Would be great if someone could share that.

upvoted 3 times

BigSteveO 4 years, 1 month ago

It's gotta be B.

upvoted 1 times

Dhiraj03 4 years, 2 months ago

Keyword : Anamolies Data prep is the only product ... So options A and C is eliminated ... Cost effective is storing the data in GCS Cloud storage ... So option is B

upvoted 1 times

nkit 4 years, 3 months ago

Dataprep to detect anomalies in Data is the right choice.

upvoted 1 times

Question #49 *Topic 1*

Google Cloud Platform resources are managed hierarchically using organization, folders, and projects. When Cloud Identity and Access Management (IAM) policies exist at these different levels, what is the effective policy at a particular node of the hierarchy?

A. The effective policy is determined only by the policy set at the node

B. The effective policy is the policy set at the node and restricted by the policies of its ancestors

C. The effective policy is the union of the policy set at the node and policies inherited from its ancestors

D. The effective policy is the intersection of the policy set at the node and policies inherited from its ancestors

Suggested Answer: *C*

Reference: https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy

Community vote distribution

C (100%)

passnow Highly Voted 6 years, 8 months ago

The effective policy for a resource is the union of the policy set at that resource and the policy inherited from its parent.https://cloud.google.com/iam/docs/resource-hierarchy-access-control

upvoted 35 times

ghadxx Highly Voted 4 years, 6 months ago

You can set IAM policies at the level of the node, in addition to policies inherited from its parent. Hence, it is a union.

upvoted 14 times

Ekramy_Elnaggar Most Recent 1 year, 9 months ago

Here's how IAM policies work in GCP's hierarchical structure: 1. Hierarchy: GCP resources are organized in a hierarchy: - Organization: The root node representing your company. - Folders: Used to organize projects within the organization. - Projects: Containers for your resources (VMs, databases, etc.). 2. Inheritance: IAM policies are inherited down the hierarchy. This means a policy set at the Organization level applies to all folders and projects within it. 3. Union of Policies: When you have policies at different levels, the effective policy at a particular node (e.g., a project) is the combination (union) of: - The policy set directly at that node. - All the policies inherited from its parent folder and the organization. Example: If a user has "Viewer" access at the Organization level and "Editor" access at the Project level, their effective permission on that project is "Editor" (the higher permission).

upvoted 5 times

Di4sa 3 years ago

From google doc: Google Cloud resources are organized hierarchically, where the organization node is the root node in the hierarchy, the projects are the children of the organization, and the other resources are descendants of projects. You can set allow policies at different levels of the resource hierarchy. Resources inherit the allow policies of the parent resource. The effective allow policy for a resource is the union of the allow policy set at that resource and the allow policy inherited from its parent.

upvoted 3 times

omermahgoub 3 years, 8 months ago

The effective policy at a particular node in the resource hierarchy in GCP is determined by the intersection of the policy set at the node and policies inherited from its ancestors, as described in option D

Cloud IAM policies in GCP are hierarchical, meaning that policies set at higher levels of the resource hierarchy can be inherited by lower levels. When a user or service account attempts to access a resource, the effective policy at that resource is determined by evaluating the policies set at the resource itself and all of its ancestors in the hierarchy. If any of the policies deny access, the user or service account will be denied access.

For example, consider the following resource hierarchy: Organization => Folder => Project => Compute Engine instance If an IAM policy is set at the organization level that allows read access to all Compute Engine instances, and a policy is set at the project level that denies read access to a specific Compute Engine instance, the effective policy for that instance will be the intersection of the two policies, which will be to deny read access to the instance.

upvoted 1 times

omermahgoub 3 years, 8 months ago

The effective policy at a particular node in the resource hierarchy in GCP is determined by the intersection of the policy set at the node and policies inherited from its ancestors, as described in option D Cloud IAM policies in GCP are hierarchical, meaning that policies set at higher levels of the

resource hierarchy can be inherited by lower levels. When a user or service account attempts to access a resource, the effective policy at that resource is determined by evaluating the policies set at the resource itself and all of its ancestors in the hierarchy. If any of the policies deny access, the user or service account will be denied access. For example, consider the following resource hierarchy: Organization => Folder => Project => Compute Engine instance If an IAM policy is set at the organization level that allows read access to all Compute Engine instances, and a policy is set at the project level that denies read access to a specific Compute Engine instance, the effective policy for that instance will be the intersection of the two policies, which will be to deny read access to the instance.Option A: The effective policy is not determined only by the policy set at the node, as policies set at higher levels in the hierarchy can also have an impact on the effective policy. Option B: The effective policy is not restricted by the policies of its ancestors, as the policies of its ancestors can also be included in the effective policy if they allow access. Option C: The effective policy is not the union of the policy set at the node and policies inherited from its ancestors, as the intersection of the policies is used to determine the effective policy.

upvoted 1 times

omermahgoub 3 years, 8 months ago

Option A: The effective policy is not determined only by the policy set at the node, as policies set at higher levels in the hierarchy can also have an impact on the effective policy. Option B: The effective policy is not restricted by the policies of its ancestors, as the policies of its ancestors can also be included in the effective policy if they allow access. Option C: The effective policy is not the union of the policy set at the node and policies inherited from its ancestors, as the intersection of the policies is used to determine the effective policy.

upvoted 1 times

habros 3 years, 8 months ago

C. Is a skewed wording question. Cannot be comprehended right away.

upvoted 2 times

megumin 3 years, 9 months ago

ok for C

upvoted 1 times

Mahmoud_E 3 years, 10 months ago

C is correct answer

upvoted 1 times

zr79 3 years, 10 months ago

English as a second language will struggle here. Good luck to us

upvoted 5 times

BiddlyBdoyng 3 years, 10 months ago

A: Would mean polcies set at the project or higher meant nothing, this is obviously wrong B: would mean you could not grant a permissions to a single VM, it would need to be at project or above (you restrict by not giving the permission) C : The permission is the sum of all the permissions you are given through the hierarchy, this is correct, you cannot restrict once it is given at a higher level. D: Would mean you would need the permission set at ancestor and the node, this would mean to get access to a single VM you would need to be given access to all VMs at the project level.

upvoted 3 times

holerina 3 years, 11 months ago

C is correct answer as it inheritance is the basic model of IAM

upvoted 2 times

avinashvidyarthi 4 years, 3 months ago

C is correct

upvoted 1 times

Atnafu 4 years, 7 months ago

C Google Cloud resources are organized hierarchically, where the organization node is the root node in the hierarchy, the projects are the children of the organization, and the other resources are descendants of projects. You can set Identity and Access Management (IAM) policies at different levels of the resource hierarchy. Resources inherit the policies of the parent resource. The effective policy for a resource is the union of the policy set at that resource and the policy inherited from its parent.

upvoted 3 times

vincy2202 4 years, 7 months ago

C is the correct answer

upvoted 2 times

haroldbenites 4 years, 8 months ago

Go for C.

upvoted 1 times

MamthaSJ 5 years, 1 month ago

Answer is C

upvoted 3 times

victory108 5 years, 3 months ago

C. The effective policy is the union of the policy set at the node and policies inherited from its ancestors

upvoted 2 times

Question #50 *Topic 1*

You are migrating your on-premises solution to Google Cloud in several phases. You will use Cloud VPN to maintain a connection between your on-premises systems and Google Cloud until the migration is completed. You want to make sure all your on-premise systems remain reachable during this period. How should you organize your networking in Google Cloud?

A. Use the same IP range on Google Cloud as you use on-premises

B. Use the same IP range on Google Cloud as you use on-premises for your primary IP range and use a secondary range that does not overlap with the range you use on-premises

C. Use an IP range on Google Cloud that does not overlap with the range you use on-premises

D. Use an IP range on Google Cloud that does not overlap with the range you use on-premises for your primary IP range and use a secondary range with the same IP range as you use on-premises

Suggested Answer: *C*
Official Answer: C

Community vote distribution



newbie2020 Highly Voted 6 years, 6 months ago
Ans is C, https://cloud.google.com/vpc/docs/using-vpc "Primary and secondary ranges can't conflict with on-premises IP ranges if you have connected your VPC network to another network with Cloud VPN, Dedicated Interconnect, or Partner Interconnect."Agreed!Perfect.. Exact find in link.agree, any ip range, shall use filewall rule to communicate, instead of setting same IP range, which is a mess to control.from the above link - it clearly states - "Primary and secondary ranges for subnets cannot overlap with any allocated range, any primary or secondary range of another subnet in the same network, or any IP ranges of subnets in peered networks." once we create a VPN, they all are part of the same network . Hence option C is correct
upvoted 133 times

newbie2020 Highly Voted 6 years, 6 months ago
Ans is C, https://cloud.google.com/vpc/docs/using-vpc

"Primary and secondary ranges can't conflict with on-premises IP ranges if you have connected your VPC network to another network with Cloud VPN, Dedicated Interconnect, or Partner Interconnect."
upvoted 133 times

Smart 6 years, 6 months ago
Agreed!
upvoted 2 times

AD2AD4 6 years, 2 months ago
Perfect.. Exact find in link.
upvoted 2 times

elaineshi 4 years, 2 months ago
agree, any ip range, shall use filewall rule to communicate, instead of setting same IP range, which is a mess to control.
upvoted 2 times

Sundeepk 6 years, 2 months ago
from the above link - it clearly states - "Primary and secondary ranges for subnets cannot overlap with any allocated range, any primary or secondary range of another subnet in the same network, or any IP ranges of subnets in peered networks." once we create a VPN, they all are part of the same network . Hence option C is correct
upvoted 13 times

KouShikyou 6 years, 10 months ago
I think C is correct.
upvoted 21 times

KouShikyou 6 years, 10 months ago
I think C is correct.Agree with C. Secondary IP range still can't overlap".... and Google Cloud until the migration is completed." Taking this as the key, the intention is to remove the connection b/w on-prem and GCP once the migration is done. and then the secondary IPs will act as primary. So I will

choose DC is okB, The key points here: - migrating in several phases - use Cloud VPN until the migration is completed - all your on-premise systems remain reachable during this periodhow to manage the routing table in VPC if is present a subnet with the same network of vpn remote net? the correct answer is CYes C it isC, no brainer. You have on-prem VPN GCP only way this data flow to work in non-over lapping subnets. You can stretch subnets at layer 7 but you wont be able to route it via VPN.

upvoted 21 times

JoeShmoe 6 years, 9 months ago

Agree with C. Secondary IP range still can't overlap".... and Google Cloud until the migration is completed." Taking this as the key, the intention is to remove the connection b/w on-prem and GCP once the migration is done. and then the secondary IPs will act as primary. So I will choose DC is okB, The key points here: - migrating in several phases - use Cloud VPN until the migration is completed - all your on-premise systems remain reachable during this periodhow to manage the routing table in VPC if is present a subnet with the same network of vpn remote net? the correct answer is C

upvoted 10 times

JoeShmoe 6 years, 9 months ago

Agree with C. Secondary IP range still can't overlap

upvoted 10 times

AWS56 6 years, 7 months ago

".... and Google Cloud until the migration is completed." Taking this as the key, the intention is to remove the connection b/w on-prem and GCP once the migration is done. and then the secondary IPs will act as primary. So I will choose D

upvoted 3 times

AWS56 6 years, 7 months ago

".... and Google Cloud until the migration is completed." Taking this as the key, the intention is to remove the connection b/w on-prem and GCP once the migration is done. and then the secondary IPs will act as primary. So I will choose DC is okB, The key points here: - migrating in several phases - use Cloud VPN until the migration is completed - all your on-premise systems remain reachable during this periodhow to manage the routing table in VPC if is present a subnet with the same network of vpn remote net? the correct answer is C

upvoted 3 times

tartar 6 years ago

C is okB, The key points here: - migrating in several phases - use Cloud VPN until the migration is completed - all your on-premise systems remain reachable during this period

upvoted 10 times

tartar 6 years ago

C is ok

upvoted 10 times

MaxNRG 4 years, 9 months ago

B, The key points here: - migrating in several phases - use Cloud VPN until the migration is completed - all your on-premise systems remain reachable during this period

upvoted 2 times

zanfo 4 years, 11 months ago

how to manage the routing table in VPC if is present a subnet with the same network of vpn remote net? the correct answer is C

upvoted 1 times

kumarp6 5 years, 9 months ago

Yes C it is

upvoted 2 times

nitinz 5 years, 5 months ago

C, no brainer. You have on-prem VPN GCP only way this data flow to work in non-over lapping subnets. You can stretch subnets at layer 7 but you wont be able to route it via VPN.

upvoted 4 times

halifax Most Recent 1 year, 5 months ago

The issue with using the same IP ranges as on-prem is purely a routing issue. The VPN tunnel secures the traffic between the networks(on-prem and GCP). However, the routing within each network still depends on unique IP addresses.

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

1. IP Address Conflicts: When you have overlapping IP ranges between your on-premises network and your Google Cloud network, you'll run into routing conflicts. Devices won't know where to send traffic, leading to connectivity problems and unreachable systems. 2. Cloud VPN and Routing: Cloud VPN creates a secure tunnel between your on-premises network and your Google Cloud Virtual Private Cloud (VPC). To ensure proper routing, each side of the connection needs to have distinct, non-overlapping IP address spaces. 3. Best Practice for Hybrid Networks: Using different IP

ranges is a standard best practice for hybrid cloud setups. It prevents ambiguity and ensures that traffic flows correctly between your on-premises and cloud environments.

upvoted 3 times

gracjanborowiak 2 years, 1 month ago

question is tricky. as network architect knowing gcp i have exp that you can use non-overlapping secondary ranges for vpn as well. in many migrations it is not possible to make new addressing hence you need to make them overlapping. this is why 2nd ranges are so useful. B is better choice. more realistic and possible in gcp. from overall perspective i agree to have non-overlapping but do not forget this is migration and you need to have full connectivity all the time. it is also not mentioning about what ips should be usedWhen migrating to the cloud, best practices for IP schema generally involve avoiding duplicate IP addresses and keeping cloud and on-premise IP ranges separate

upvoted 1 times

gracjanborowiak 2 years, 1 month ago

Selected Answer: B

question is tricky. as network architect knowing gcp i have exp that you can use non-overlapping secondary ranges for vpn as well.

in many migrations it is not possible to make new addressing hence you need to make them overlapping. this is why 2nd ranges are so useful.

B is better choice. more realistic and possible in gcp.

from overall perspective i agree to have non-overlapping but do not forget this is migration and you need to have full connectivity all the time. it is also not mentioning about what ips should be used

upvoted 1 times

desertlotus1211 2 years ago

When migrating to the cloud, best practices for IP schema generally involve avoiding duplicate IP addresses and keeping cloud and on-premise IP ranges separate

upvoted 1 times

heretolearnazure 2 years, 12 months ago

C is correct

upvoted 1 times

JC0926 3 years, 5 months ago

Using an IP range on Google Cloud that does not overlap with the range used on-premises (option C) is a good choice to avoid IP address conflicts. However, it is important to use the same IP range as the on-premises applications for the primary IP range to ensure that the on-premises systems remain accessible. Therefore, using the same IP range on Google Cloud as on-premises for the primary IP range and using a secondary range that does not overlap with the range used on-premises can avoid IP address duplication and ensure that the on-premises systems remain accessible. Hence, option B is the better choice.

upvoted 3 times

omermahgoub 3 years, 8 months ago

The recommended approach for organizing your networking in Google Cloud to ensure that all your on-premises systems remain reachable during the migration is option C: Use an IP range on Google Cloud that does not overlap with the range you use on-premises. When using Cloud VPN to establish a connection between your on-premises systems and Google Cloud, it is important to ensure that the IP ranges used in your on-premises systems and Google Cloud do not overlap. If the IP ranges overlap, it can cause conflicts and make it difficult to route traffic between your on-premises systems and Google Cloud. To avoid IP range conflicts, you should use an IP range on Google Cloud that is different from the range you use on-premises. This will ensure that all your on-premises systems remain reachable during the migration.Option A: Using the same IP range on Google Cloud as you use on-premises is not a recommended approach, as it can cause IP range conflicts and make it difficult to route traffic between your on-premises systems and Google Cloud. Option B: Using the same IP range on Google Cloud as you use on-premises for your primary IP range and a secondary range that does not overlap with the range you use on-premises is not a recommended approach, as it can still cause IP range conflicts and make it difficult to route traffic between your on-premises systems and Google Cloud. Option D: Using an IP range on Google Cloud that does not overlap with the range you use on-premises for your primary

upvoted 2 times

omermahgoub 3 years, 8 months ago

The recommended approach for organizing your networking in Google Cloud to ensure that all your on-premises systems remain reachable during the migration is option C: Use an IP range on Google Cloud that does not overlap with the range you use on-premises.

When using Cloud VPN to establish a connection between your on-premises systems and Google Cloud, it is important to ensure that the IP ranges used in your on-premises systems and Google Cloud do not overlap. If the IP ranges overlap, it can cause conflicts and make it difficult to route traffic between your on-premises systems and Google Cloud.

To avoid IP range conflicts, you should use an IP range on Google Cloud that is different from the range you use on-premises. This will ensure that all your on-premises systems remain reachable during the migration.

upvoted 2 times

omermahgoub 3 years, 8 months ago

Option A: Using the same IP range on Google Cloud as you use on-premises is not a recommended approach, as it can cause IP range conflicts and make it difficult to route traffic between your on-premises systems and Google Cloud. Option B: Using the same IP range on Google Cloud as you use on-premises for your primary IP range and a secondary range that does not overlap with the range you use on-premises is not a recommended approach, as it can still cause IP range conflicts and make it difficult to route traffic between your on-premises systems and Google Cloud. Option D: Using an IP range on Google Cloud that does not overlap with the range you use on-premises for your primary

upvoted 1 times

megumin 3 years, 9 months ago

ok for C

upvoted 1 times

zr79 3 years, 10 months ago

no overlapping

upvoted 1 times

AzureDP900 3 years, 10 months ago

C. Use an IP range on Google Cloud that does not overlap with the range you use on-premises

upvoted 1 times

marksie1988 3 years, 11 months ago

C, IP should never overlap if avoidable. double nat is nasty

upvoted 1 times

ZLT 4 years, 1 month ago

The correct answer is C

upvoted 2 times

Barry123456 4 years, 2 months ago

C Why would you ever create an IP overlap?

upvoted 1 times

jonty4gcp 4 years, 4 months ago

Answer is C

upvoted 1 times

Davidik79 4 years, 5 months ago

From here: https://cloud.google.com/vpc/docs/create-modify-vpc-networks "Primary and secondary ranges can't conflict with on-premises IP ranges if you have connected your VPC network to another network with Cloud VPN, Dedicated Interconnect, or Partner Interconnect."

upvoted 1 times

Sreedharveluru 4 years, 7 months ago

ANS - C Primary and secondary ranges for subnets cannot overlap with any allocated range, any primary or secondary range of another subnet in the same network, or any IPv4 ranges of subnets in peered networks.

upvoted 2 times

Question #51 *Topic 1*

You have found an error in your App Engine application caused by missing Cloud Datastore indexes. You have created a YAML file with the required indexes and want to deploy these new indexes to Cloud Datastore. What should you do?

A. Point gcloud datastore create-indexes to your configuration file

B. Upload the configuration file to App Engine's default Cloud Storage bucket, and have App Engine detect the new indexes

C. In the GCP Console, use Datastore Admin to delete the current indexes and upload the new configuration file

D. Create an HTTP request to the built-in python module to send the index configuration file to your application

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (100%)

jcmoranp Highly Voted 6 years, 9 months ago
Correct A, you have to recreate the indexesA, if index is missing then create it.A is okYes, use this command in cloud sell to create indexes. gcloud datastore create indexes
upvoted 31 times

jcmoranp Highly Voted 6 years, 9 months ago
Correct A, you have to recreate the indexes
upvoted 31 times

nitinz 5 years, 5 months ago
A, if index is missing then create it.
upvoted 4 times

tartar 6 years ago
A is ok
upvoted 11 times

kumarp6 5 years, 9 months ago
Yes, use this command in cloud sell to create indexes. gcloud datastore create indexes
upvoted 4 times

Eroc 6 years, 9 months ago
A is incorrect because the command is actually gcloud datastore indexes create. (https://cloud.google.com/sdk/gcloud/reference/datastore/indexes/create).
upvoted 17 times

Eroc 6 years, 9 months ago
A is incorrect because the command is actually gcloud datastore indexes create. (https://cloud.google.com/sdk/gcloud/reference/datastore/indexes/create).It might have changed recently - I was able to find documentation mentioning "datastore create-indexes": https://cloud.google.com/appengine/docs/standard/python/datastore/indexes
upvoted 17 times

bogd 5 years, 6 months ago
It might have changed recently - I was able to find documentation mentioning "datastore create-indexes": https://cloud.google.com/appengine/docs/standard/python/datastore/indexes
upvoted 15 times

sim7243 Most Recent 1 year, 9 months ago
A, option
upvoted 1 times

james2033 2 years, 2 months ago
To deploy new Cloud Datastore indexes using the YAML file you created, you should use the `gcloud` command-line tool. Specifically, the correct option is: **A. Point gcloud datastore create-indexes to your configuration file** Here's how you can do it: 1. **Ensure you have the `gcloud` CLI installed**: You need the Google Cloud SDK installed and set up on your local machine. If you haven't done this yet, follow the [installation guide](https://cloud.google.com/sdk/docs/install). 2. **Navigate to the directory containing your `index.yaml` file**: This file contains the definitions of the

indexes you want to deploy. 3. **Run the following command**: ```bash gcloud datastore indexes create index.yaml ``` This command will deploy the indexes defined in the `index.yaml` file to Cloud Datastore.

upvoted 1 times

thewalker 2 years, 9 months ago

A is correct: https://cloud.google.com/sdk/gcloud/reference/datastore/indexes/create

upvoted 3 times

jlambdan 3 years, 5 months ago

A) https://cloud.google.com/datastore/docs/tools/indexconfig#Datastore_Updating_indexes

upvoted 2 times

MestreCholas 3 years, 5 months ago

A. Point gcloud datastore create-indexes to your configuration file. To deploy new indexes to Cloud Datastore, you can use the gcloud datastore create-indexes command and point it to the YAML configuration file containing the required indexes. This command will create the new indexes in Cloud Datastore for your application. Option B is not correct because App Engine does not automatically detect and create indexes from uploaded configuration files in Cloud Storage. Option C is also not correct because deleting current indexes in Datastore Admin is not necessary to upload new indexes. Option D is not correct because there is no built-in Python module that can send the index configuration file to your application

upvoted 12 times

omermahgoub 3 years, 8 months ago

To deploy new indexes to Cloud Datastore, you should use the gcloud datastore create-indexes command and point it to your configuration file. The correct option is therefore A: Point gcloud datastore create-indexes to your configuration file. Here's the general format of the gcloud datastore create-indexes command: Copy code gcloud datastore create-indexes [FILE] Where [FILE] is the path to your configuration file. The configuration file should be in YAML format and contain a list of indexes that you want to create. For example: Copy code gcloud datastore create-indexes index.yaml This command will create the indexes specified in the index.yaml file in Cloud Datastore.

upvoted 7 times

megumin 3 years, 9 months ago

ok for A

upvoted 1 times

Mahmoud_E 3 years, 10 months ago

A is correct answer

upvoted 1 times

AzureDP900 3 years, 10 months ago

A is correct

upvoted 1 times

minmin2020 3 years, 10 months ago

gcloud app deploy index.yaml

upvoted 2 times

AMEJack 3 years, 10 months ago

Answer is A. Option C: no create index from configuration file in the firestore console.

upvoted 1 times

Najmuddoja 3 years, 10 months ago

Is A still the valid answer, many other website resources saying C is the right answer

upvoted 1 times

holerina 3 years, 11 months ago

A looks like more logical answer

upvoted 1 times

vincy2202 4 years, 7 months ago

A is the correct answer

upvoted 1 times

menon_sarath 4 years, 8 months ago

Why not C? Shouldnt the command be datastore indexes create and not create-index as incorrectly stated in option A? Please advice

upvoted 3 times

Question #52 *Topic 1*

You have an application that will run on Compute Engine. You need to design an architecture that takes into account a disaster recovery plan that requires your application to fail over to another region in case of a regional outage. What should you do?

A. Deploy the application on two Compute Engine instances in the same project but in a different region. Use the first instance to serve traffic, and use the HTTP load balancing service to fail over to the standby instance in case of a disaster.

B. Deploy the application on a Compute Engine instance. Use the instance to serve traffic, and use the HTTP load balancing service to fail over to an instance on your premises in case of a disaster.

C. Deploy the application on two Compute Engine instance groups, each in the same project but in a different region. Use the first instance group to serve traffic, and use the HTTP load balancing service to fail over to the standby instance group in case of a disaster.

D. Deploy the application on two Compute Engine instance groups, each in a separate project and a different region. Use the first instance group to serve traffic, and use the HTTP load balancing service to fail over to the standby instance group in case of a disaster.

Suggested Answer: *C*
Official Answer: C

Community vote distribution
C (100%)

Eroc Highly Voted 6 years, 3 months ago
Groups are better for management that non-groups so A and B are eliminated. Keeping the the instances in the same project will help maintain consistency, so C is better than D.
upvoted 39 times

Eroc Highly Voted 6 years, 3 months ago
Groups are better for management that non-groups so A and B are eliminated. Keeping the the instances in the same project will help maintain consistency, so C is better than D.make sense.C, because external LB needs **IG** period. It can either be managed or un-managed. You can not do External HTTP LB on instances. Also External HHTP LB is a Regional resource.
upvoted 39 times

LaxmanTiwari 2 years, 8 months ago
make sense.
upvoted 1 times

nitinz 4 years, 11 months ago
C, because external LB needs **IG** period. It can either be managed or un-managed. You can not do External HTTP LB on instances. Also External HHTP LB is a Regional resource.
upvoted 15 times

gigibit 4 years, 5 months ago
Yes, but why not choose a cost-effective solution like A, preferring a not required performance optimization solution like C? The question it's just asking for a simple fail over
upvoted 13 times

gigibit 4 years, 5 months ago
Yes, but why not choose a cost-effective solution like A, preferring a not required performance optimization solution like C? The question it's just asking for a simple fail overLB does not work without MIGUse two groups, each group contain one instance would be more flixible than answer A.In real word situation, I'd choose A for a customer. The question doesn't mention performance to be enhanced or scalability.
upvoted 13 times

hellosam 2 years, 6 months ago
LB does not work without MIG
upvoted 4 times

Terryhsieh Most Recent 2 years, 1 month ago
Use two groups, each group contain one instance would be more flixible than answer A.
upvoted 2 times

todos213 3 years, 10 months ago
In real word situation, I'd choose A for a customer. The question doesn't mention performance to be enhanced or scalability.

upvoted 6 times

yilexar 2 years, 4 months ago

Now you can use regional load balancer to route traffic to instance groups in different GCP projects. Welcome to cloud :-)
https://cloud.google.com/blog/products/networking/cloud-load-balancing-gets-cross-project-service-referencing

upvoted 6 times

heretolearnazure 2 years, 5 months ago

C is correct

upvoted 1 times

BeCalm 2 years, 11 months ago

To set up a load balancer with a Compute Engine backend, your VMs need to be in an instance group. The managed instance group provides VMs running the backend servers of an external HTTP load balancer

Therefore C

upvoted 6 times

BeCalm 2 years, 11 months ago

To set up a load balancer with a Compute Engine backend, your VMs need to be in an instance group. The managed instance group provides VMs running the backend servers of an external HTTP load balancer Therefore Cgood catch .

upvoted 6 times

LaxmanTiwari 2 years, 8 months ago

good catch .

upvoted 2 times

BeCalm 2 years, 11 months ago

Why can't this be A since there is no mention of scaling?

upvoted 1 times

BeCalm 2 years, 11 months ago

Why can't this be A since there is no mention of scaling?LB does not work without MIGYou can use either MIG or unmanaged..
https://cloud.google.com/compute/docs/instance-groups/creating-groups-of-unmanaged-instances

upvoted 1 times

hellosam 2 years, 6 months ago

LB does not work without MIGYou can use either MIG or unmanaged.. https://cloud.google.com/compute/docs/instance-groups/creating-groups-of-unmanaged-instances

upvoted 1 times

axle818 2 years, 2 months ago

You can use either MIG or unmanaged.. https://cloud.google.com/compute/docs/instance-groups/creating-groups-of-unmanaged-instances

upvoted 2 times

n_nana 3 years, 1 month ago

Google recommend using MIG for Zonal outage and multiple MIG for regional outage https://cloud.google.com/architecture/disaster-recovery#compute-engine sentence says: Compute Engine instances are zonal resources, so in the event of a zone outage instances are unavailable by default. Compute Engine does offer managed instance groups (MIGs) which can automatically scale up additional VMs from pre-configured instance templates, both within a single zone and across multiple zones within a region. MIGs are ideal for applications that require resilience to zone loss and are stateless, but require configuration and resource planning. Multiple regional MIGs can be used to achieve region outage resilience for stateless applications.

upvoted 8 times

roaming_panda 3 years, 1 month ago

MIG , in 1 project

upvoted 1 times

sameer2803 3 years, 1 month ago

the only solution that they want is to address regional outage. scalability or performance is not a concern at this point. MIG is better but that not what is solving the regional outage.According to this link, MIG is not only for performance and scalability. It is also for reliability
https://cloud.google.com/architecture/disaster-recovery#compute-engine sentence says: Compute Engine instances are zonal resources, so in the event of a zone outage instances are unavailable by default. Compute Engine does offer managed instance groups (MIGs) which can automatically scale up additional VMs from pre-configured instance templates, both within a single zone and across multiple zones within a region. MIGs are ideal for applications that require resilience to zone loss and are stateless, but require configuration and resource planning. Multiple regional MIGs can be used to achieve region outage resilience for stateless applications.

upvoted 1 times

sameer2803 3 years, 1 month ago

the only solution that they want is to address regional outage. scalability or performance is not a concern at this point. MIG is better but that not what is solving the regional outage.

upvoted 1 times

n_nana 3 years, 1 month ago

According to this link, MIG is not only for performance and scalability. It is also for reliability https://cloud.google.com/architecture/disaster-recovery#compute-engine sentence says: Compute Engine instances are zonal resources, so in the event of a zone outage instances are unavailable by default. Compute Engine does offer managed instance groups (MIGs) which can automatically scale up additional VMs from pre-configured instance templates, both within a single zone and across multiple zones within a region. MIGs are ideal for applications that require resilience to zone loss and are stateless, but require configuration and resource planning. Multiple regional MIGs can be used to achieve region outage resilience for stateless applications.

upvoted 1 times

omermahgoub 3 years, 1 month ago

The correct answer is D: Deploy the application on two Compute Engine instance groups, each in a separate project and a different region. Use the first instance group to serve traffic, and use the HTTP load balancing service to fail over to the standby instance group in case of a disaster.

To implement a disaster recovery plan that requires your application to fail over to another region in case of a regional outage, you should deploy the application on two Compute Engine instance groups, each in a separate project and a different region. This will ensure that the application is running in at least two regions, so that if one region experiences an outage, the application can still be accessed from the other region.

You can use the HTTP load balancing service to distribute traffic between the two instance groups and to fail over to the standby instance group in case of a disaster. This will ensure that your application is always available, even in the event of a regional outage.

upvoted 2 times

omermahgoub 3 years, 1 month ago

The correct answer is D: Deploy the application on two Compute Engine instance groups, each in a separate project and a different region. Use the first instance group to serve traffic, and use the HTTP load balancing service to fail over to the standby instance group in case of a disaster. To implement a disaster recovery plan that requires your application to fail over to another region in case of a regional outage, you should deploy the application on two Compute Engine instance groups, each in a separate project and a different region. This will ensure that the application is running in at least two regions, so that if one region experiences an outage, the application can still be accessed from the other region. You can use the HTTP load balancing service to distribute traffic between the two instance groups and to fail over to the standby instance group in case of a disaster. This will ensure that your application is always available, even in the event of a regional outage.Option A: Deploying the application on two Compute Engine instances in the same project but in a different region will not provide enough redundancy, as the instances are still in the same project and could be affected by the same regional outage. Option B: Deploying the application on a Compute Engine instance and using the HTTP load balancing service to fail over to an instance on your premises in case of a disaster is not a valid option, as it does not provide the required disaster recovery capability. Option C: Deploying the application on two Compute Engine instance groups, each in the same project but in a different region, is not a valid option, as it does not provide the required disaster recovery capability.

upvoted 2 times

omermahgoub 3 years, 1 month ago

Option A: Deploying the application on two Compute Engine instances in the same project but in a different region will not provide enough redundancy, as the instances are still in the same project and could be affected by the same regional outage. Option B: Deploying the application on a Compute Engine instance and using the HTTP load balancing service to fail over to an instance on your premises in case of a disaster is not a valid option, as it does not provide the required disaster recovery capability. Option C: Deploying the application on two Compute Engine instance groups, each in the same project but in a different region, is not a valid option, as it does not provide the required disaster recovery capability.

upvoted 2 times

oms_muc 3 years, 2 months ago

IMHO A would also work (test env). In production relying on instances from other regions (introducing latency) just for possible expected zonal outages, would not be my best practice.

upvoted 1 times

AniketD 3 years, 3 months ago

Correct answer is C. No need to host MIG in separate projects.

upvoted 2 times

megumin 3 years, 3 months ago

ok for C

upvoted 1 times

AzureDP900 3 years, 4 months ago

C is right choice

upvoted 1 times

minmin2020 3 years, 4 months ago

C - using instance groups when planning for DR is better than having single vm's (https://cloud.google.com/architecture/disaster-recovery). Having the resources in the same project is probably good for resource management.

upvoted 1 times

arjunvijayvargiya 3 years, 4 months ago

Answer should have been A as autoscaling requirements are not mentioned.

upvoted 1 times

BiddlyBdoyng 3 years, 4 months ago

I think you need an group for the load balancer so A no good?

upvoted 2 times

Question #53 *Topic 1*

You are deploying an application on App Engine that needs to integrate with an on-premises database. For security purposes, your on-premises database must not be accessible through the public internet. What should you do?

A. Deploy your application on App Engine standard environment and use App Engine firewall rules to limit access to the open on-premises database.

B. Deploy your application on App Engine standard environment and use Cloud VPN to limit access to the on-premises database.

C. Deploy your application on App Engine flexible environment and use App Engine firewall rules to limit access to the on-premises database.

D. Deploy your application on App Engine flexible environment and use Cloud VPN to limit access to the on-premises database.

Suggested Answer: *D*

Official Answer: D

Community vote distribution

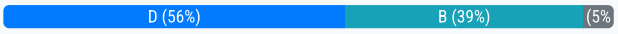

MyPractice Highly Voted 6 years, 7 months ago

Agree with D - "When to choose the flexible environment" "Accesses the resources or services of your Google Cloud project that reside in the Compute Engine network." https://cloud.google.com/appengine/docs/the-appengine-environments

upvoted 54 times

MyPractice Highly Voted 6 years, 7 months ago

Agree with D - "When to choose the flexible environment" "Accesses the resources or services of your Google Cloud project that reside in the Compute Engine network." https://cloud.google.com/appengine/docs/the-appengine-environmentsWhy not B ? https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databasesIsn't the question said "not public internet access"?Yes, that phrase in the question bothers me too. However, when I check this: https://cloud.google.com/appengine/docs/flexible/storage-options#:~:text=On%20premises,-If%20you%20have&text=Because%20App%20Engine%20and%20Compute,database%20server's%20internal%20IP%20address. it says "If you have existing on-premises databases that you want to make accessible to your App Engine app, you can either configure your internal network and firewall to give the database a public IP address or connect using a VPN." So I think the question should have skipped the words "not public internet access" if they want us to choose VPN.In a forum mentions that GCE and CAP flex are designed for connect to VPC . With GAP standard is needed a proxy . https://stackoverflow.com/questions/47537204/how-to-connect-app-engine-and-on-premise-server-through-vpnbecause app engine standard cant connect to on-prem dbWhere does it say appengine cannot connect to on-prem db? With CloudVPN, it shoudl connect as per this https://cloud.google.com/appengine/docs/flexible/storage-options#on_premises Also going with D will require app to be containerized. That is not listed in the requirement.Your link points to a capability of Flexible environment, not Standard.

upvoted 54 times

AWS56 6 years, 7 months ago

Why not B ? https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databasesIsn't the question said "not public internet access"?Yes, that phrase in the question bothers me too. However, when I check this: https://cloud.google.com/appengine/docs/flexible/storage-options#:~:text=On%20premises,-If%20you%20have&text=Because%20App%20Engine%20and%20Compute,database%20server's%20internal%20IP%20address. it says "If you have existing on-premises databases that you want to make accessible to your App Engine app, you can either configure your internal network and firewall to give the database a public IP address or connect using a VPN." So I think the question should have skipped the words "not public internet access" if they want us to choose VPN.In a forum mentions that GCE and CAP flex are designed for connect to VPC . With GAP standard is needed a proxy . https://stackoverflow.com/questions/47537204/how-to-connect-app-engine-and-on-premise-server-through-vpnbecause app engine standard cant connect to on-prem dbWhere does it say appengine cannot connect to on-prem db? With CloudVPN, it shoudl connect as per this https://cloud.google.com/appengine/docs/flexible/storage-options#on_premises Also going with D will require app to be containerized. That is not listed in the requirement.Your link points to a capability of Flexible environment, not Standard.

upvoted 7 times

AWS56 6 years, 7 months ago

Why not B ? https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databases

upvoted 7 times

elaineshi 4 years, 2 months ago

Isn't the question said "not public internet access"?Yes, that phrase in the question bothers me too. However, when I check this: https://cloud.google.com/appengine/docs/flexible/storage-options#:~:text=On%20premises,-If%20you%20have&text=Because%20App%20Engine%20and%20Compute,database%20server's%20internal%20IP%20address. it says "If you have existing on-premises databases that you want to make accessible to your App Engine app, you can either configure your internal network and firewall to give the database a public IP address or connect using a VPN." So I think the question should have skipped the words "not public internet access" if they want us to choose VPN.

upvoted 2 times

elaineshi 4 years, 2 months ago

Isn't the question said "not public internet access"?

upvoted 2 times

mnsait 2 years, 3 months ago

Yes, that phrase in the question bothers me too. However, when I check this: https://cloud.google.com/appengine/docs/flexible/storage-options#:~:text=On%20premises,-If%20you%20have&text=Because%20App%20Engine%20and%20Compute,database%20server's%20internal%20IP%20address. it says "If you have existing on-premises databases that you want to make accessible to your App Engine app, you can either configure your internal network and firewall to give the database a public IP address or connect using a VPN." So I think the question should have skipped the words "not public internet access" if they want us to choose VPN.

upvoted 3 times

haroldbenites 4 years, 8 months ago

In a forum mentions that GCE and CAP flex are designed for connect to VPC . With GAP standard is needed a proxy . https://stackoverflow.com/questions/47537204/how-to-connect-app-engine-and-on-premise-server-through-vpn

upvoted 5 times

areza 5 years, 2 months ago

because app engine standard cant connect to on-prem dbWhere does it say appengine cannot connect to on-prem db? With CloudVPN, it shoudl connect as per this https://cloud.google.com/appengine/docs/flexible/storage-options#on_premises Also going with D will require app to be containerized. That is not listed in the requirement.Your link points to a capability of Flexible environment, not Standard.

upvoted 31 times

areza 5 years, 2 months ago

because app engine standard cant connect to on-prem db

upvoted 31 times

VSMu 3 years, 6 months ago

Where does it say appengine cannot connect to on-prem db? With CloudVPN, it shoudl connect as per this https://cloud.google.com/appengine/docs/flexible/storage-options#on_premises Also going with D will require app to be containerized. That is not listed in the requirement.Your link points to a capability of Flexible environment, not Standard.

upvoted 6 times

VSMu 3 years, 6 months ago

Where does it say appengine cannot connect to on-prem db? With CloudVPN, it shoudl connect as per this https://cloud.google.com/appengine/docs/flexible/storage-options#on_premises

Also going with D will require app to be containerized. That is not listed in the requirement.

upvoted 6 times

BeCalm 3 years, 5 months ago

Your link points to a capability of Flexible environment, not Standard.

upvoted 3 times

jcmoranp 6 years, 9 months ago

Right is D: https://stackoverflow.com/questions/37137914/is-it-possible-to-use-google-app-engine-with-google-cloud-vpnQuestion is can we restrict acess with VP N ?The stackoverflow reference if older that the answer (6 years) I think that has changed.

upvoted 18 times

jcmoranp 6 years, 9 months ago

Right is D: https://stackoverflow.com/questions/37137914/is-it-possible-to-use-google-app-engine-with-google-cloud-vpn

upvoted 18 times

amxexam 4 years, 3 months ago

Question is can we restrict acess with VP N ?The stackoverflow reference if older that the answer (6 years) I think that has changed.

upvoted 5 times

amxexam 4 years, 3 months ago

Question is can we restrict acess with VP N ?

upvoted 5 times

moiradavis 4 years ago

The stackoverflow reference if older that the answer (6 years) I think that has changed.

upvoted 1 times

keyik777 Most Recent 1 year, 3 months ago

VPN and Flex Env

upvoted 2 times

cloud_rider 1 year, 5 months ago

D is Correct, check this link - https://cloud.google.com/appengine/docs/flexible/flexible-for-standard-users#when_to_use_the Flexible allows third party libraries

upvoted 2 times

alihabib 1 year, 6 months ago

D - Flexible environment gives access to resources (Networking & Compute), additionally Cloud VPN gives private access to On-Premises

upvoted 3 times

deep316 1 year, 8 months ago

Standard requires more setup compared to Flexible. Standard Environment: To connect from the standard environment, you primarily use "Serverless VPC Access" which allows your App Engine app to reach your VPC network over private IP addresses without exposing it directly to the public internet. Flexible Environment: In the flexible environment, you can directly connect to your VPC network by deploying your app within the same VPC as your Cloud VPN gateway, enabling a more seamless connection using the private IP addresses of your network resources.

upvoted 4 times

Nimeshv 1 year, 8 months ago

App Engine Standard Environment: Limited Customization: It runs on predefined runtime environments with limited flexibility. No Access to Compute Engine Network: It does not support integration with the VPC network, which is necessary for setting up VPN connections. App Engine Flexible Environment: Customizable Runtimes: Allows you to run custom runtimes and use your own Docker containers. VPC Integration: Fully supports VPC networking, enabling you to set up Cloud VPN to securely connect to your on-premises database. Greater Flexibility: More control over the instance types, scaling, and networking options.

upvoted 4 times

desertlotus1211 1 year, 8 months ago

Per google: For using Cloud VPN, App Engine Flexible is generally considered better than App Engine Standard because it offers more customization and control over your virtual machine environment, allowing you to configure network settings and access on-premises resources more easily through the VPN Answer is D

upvoted 3 times

desertlotus1211 1 year, 8 months ago

Selected Answer: D

The answer is BOTH B&D...

Answer: BD I'm not sure what is the point of the question or the problem. The standard environment can scale from zero instances up to thousands very quickly. In contrast, the flexible environment must have at least one instance running for each active version and can take longer to scale out in response to traffic. Standard environment uses a custom-designed autoscaling algorithm.

But the question doesn't address this.

upvoted 3 times

desertlotus1211 1 year, 8 months ago

The answer is BOTH B&D... Answer: BD I'm not sure what is the point of the question or the problem. The standard environment can scale from zero instances up to thousands very quickly. In contrast, the flexible environment must have at least one instance running for each active version and can take longer to scale out in response to traffic. Standard environment uses a custom-designed autoscaling algorithm. But the question doesn't address this.Sorry I have to mark an answer. I wanted to choose B&D, but it won't let me

upvoted 3 times

desertlotus1211 1 year, 8 months ago

Sorry I have to mark an answer. I wanted to choose B&D, but it won't let me

upvoted 1 times

icarogsm 1 year, 9 months ago

D, just App Engine Flex can connect to onprem.

upvoted 3 times

Ekramy_Elnaggar 1 year, 9 months ago

The flexible environment gives you more control over the networking configuration of your application. This is crucial for setting up a secure connection to your on-premises database.

upvoted 4 times

yocixim836 1 year, 9 months ago

https://cloud.google.com/appengine/docs/standard/storage-options#on_premises

upvoted 3 times

dpttpd 1 year, 10 months ago

https://cloud.google.com/appengine/docs/standard/storage-options#on_premises

upvoted 3 times

raghupothula 1 year, 10 months ago

B and C are not appropriate and wherein in App Engine flex resources reside in VPC NEtwo https://cloud.google.com/appengine/docs/the-appengine-environments#app-engine-environments D

upvoted 1 times

maxdanny 1 year, 11 months ago

App Engine flexible environment provides more flexibility and supports VPC (Virtual Private Cloud) connectivity, which allows you to set up a Cloud VPN connection. The VPN can be used to securely connect your App Engine application to the on-premises database without exposing it to the public internet.

upvoted 4 times

joecloud12 2 years ago

flexible is more expensive. standard will suffice

upvoted 1 times

janji456 2 years ago

D et up a VPN connection between your on-premises network and Google Cloud. This establishes a secure tunnel for communication. Your App Engine

upvoted 1 times

Question #54 *Topic 1*

You are working in a highly secured environment where public Internet access from the Compute Engine VMs is not allowed. You do not yet have a VPN connection to access an on-premises file server. You need to install specific software on a Compute Engine instance. How should you install the software?

A. Upload the required installation files to Cloud Storage. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gsutil.

B. Upload the required installation files to Cloud Storage and use firewall rules to block all traffic except the IP address range for Cloud Storage. Download the files to the VM using gsutil.

C. Upload the required installation files to Cloud Source Repositories. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gcloud.

D. Upload the required installation files to Cloud Source Repositories and use firewall rules to block all traffic except the IP address range for Cloud Source Repositories. Download the files to the VM using gsutil.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (82%) B (18%)

zaki_b Highly Voted 6 years, 9 months ago
Internet access is not allowed so it should be A. CMIIWA is okA is the answerA is the best answer.
upvoted 55 times

zaki_b Highly Voted 6 years, 9 months ago
Internet access is not allowed so it should be A. CMIIW
upvoted 55 times

tartar 6 years ago
A is ok
upvoted 10 times

kumarp6 5 years, 9 months ago
A is the answer
upvoted 3 times

nitinz 5 years, 5 months ago
A is the best answer.
upvoted 3 times

KNG 6 years, 6 months ago
Should be A https://cloud.google.com/vpc/docs/configure-private-services-access Note: Even though the IP addresses for Google APIs and services are public, the traffic path from instances that are using Private Google Access to the Google APIs remains within Google's network.
upvoted 21 times

plumbig11 Most Recent 1 year, 7 months ago
Private Google Access subnet with cloud storage. A
upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago
Private Google Access allows your VM to access Google APIs and services (like Cloud Storage) without needing a public IP address. This is crucial in your restricted environment.
upvoted 1 times

maxdanny 1 year, 11 months ago
Private Google Access ensures the VM can reach Cloud Storage using its internal IP, while still restricting public internet access.
upvoted 3 times

arotesa 1 year, 11 months ago
The Answer is D
upvoted 1 times

AdityaGupta 2 years, 10 months ago

A. Upload the required installation files to Cloud Storage. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gsutil.

upvoted 1 times

ppandher 3 years, 7 months ago

Those who are opting for B, Can please explain without Internet access and without Private Google Access enabled how will they communicate with Cloud Storage ? :)

upvoted 5 times

omermahgoub 3 years, 8 months ago

The correct answer is A: Upload the required installation files to Cloud Storage. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gsutil.

To install specific software on a Compute Engine instance in a highly secured environment where public Internet access is not allowed, you can follow these steps:

Upload the required installation files to Cloud Storage. Configure the VM on a subnet with a Private Google Access subnet. This will allow the VM to access Google APIs and services, such as Cloud Storage, without requiring a public IP address or internet access. Assign only an internal IP address to the VM. This will ensure that the VM is not accessible from the public internet. Download the installation files to the VM using gsutil, which is a command-line tool that allows you to access Cloud Storage from the VM.

upvoted 8 times

omermahgoub 3 years, 8 months ago

The correct answer is A: Upload the required installation files to Cloud Storage. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gsutil. To install specific software on a Compute Engine instance in a highly secured environment where public Internet access is not allowed, you can follow these steps: Upload the required installation files to Cloud Storage. Configure the VM on a subnet with a Private Google Access subnet. This will allow the VM to access Google APIs and services, such as Cloud Storage, without requiring a public IP address or internet access. Assign only an internal IP address to the VM. This will ensure that the VM is not accessible from the public internet. Download the installation files to the VM using gsutil, which is a command-line tool that allows you to access Cloud Storage from the VM.Option B: Uploading the required installation files to Cloud Storage and using firewall rules to block all traffic except the IP address range for Cloud Storage is not a valid option, as it does not allow the VM to access the installation files without public internet access. Option C: Uploading the required installation files to Cloud Source Repositories and using gcloud to download the files to the VM is not a valid option, as Cloud Source Repositories does not support storing large binary files such as installation files. Option D: Uploading the required installation files to Cloud Source Repositories and using firewall rules to block all traffic except the IP address range for Cloud Source Repositories is not a valid option, as it does not allow the VM to access the installation files without public internet access.

upvoted 8 times

omermahgoub 3 years, 8 months ago

Option B: Uploading the required installation files to Cloud Storage and using firewall rules to block all traffic except the IP address range for Cloud Storage is not a valid option, as it does not allow the VM to access the installation files without public internet access. Option C: Uploading the required installation files to Cloud Source Repositories and using gcloud to download the files to the VM is not a valid option, as Cloud Source Repositories does not support storing large binary files such as installation files. Option D: Uploading the required installation files to Cloud Source Repositories and using firewall rules to block all traffic except the IP address range for Cloud Source Repositories is not a valid option, as it does not allow the VM to access the installation files without public internet access.

upvoted 6 times

habros 3 years, 8 months ago

Eliminate B&D as it connect via public networks despite it being a Google Cloud service.

upvoted 1 times

megumin 3 years, 9 months ago

ok for A

upvoted 1 times

stevehlw 3 years, 9 months ago

With private Google access subnet, the vm can reach external network. With this setting, it violates “public Internet access from the Compute Engine VMs is not allowed”. Can someone explain why it’s not B instead?

upvoted 2 times

stevehlw 3 years, 9 months ago

With private Google access subnet, the vm can reach external network. With this setting, it violates “public Internet access from the Compute Engine VMs is not allowed”. Can someone explain why it’s not B instead?Private Google access means - refer to https://www.youtube.com/watch?

v=yd5FtV8aJkk

upvoted 2 times

ppandher 3 years, 8 months ago

Private Google access means - refer to https://www.youtube.com/watch?v=yd5FtV8aJkk

upvoted 3 times

AzureDP900 3 years, 10 months ago

A is good

upvoted 1 times

minmin2020 3 years, 10 months ago

A. Upload the required installation files to Cloud Storage. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gsutil.

upvoted 1 times

muneebarshad 3 years, 11 months ago

Configuring Private Google Access is the best way to access Google Services for VM that does not have access to the internet. In order to access Google Private APIs egress should be opened to the following IP Address restricted.googleapis.com (199.36.153.4/30). VM will leverage internal networking to access Cloud Storage https://cloud.google.com/vpc/docs/configure-private-google-access

upvoted 4 times

6721sora 3 years, 12 months ago

C because Cloud repositories is a private Git within Google cloud. Hence it is ideal for simple pull, push, clone type "git" operations. As this is within Google cloud and is a private git, you do not need public internet accessI think it's not this because Clouse Source Repositories is for source code. Sounds like we are looking for an executable?

upvoted 1 times

6721sora 3 years, 12 months ago

C because Cloud repositories is a private Git within Google cloud. Hence it is ideal for simple pull, push, clone type "git" operations. As this is within Google cloud and is a private git, you do not need public internet access

upvoted 1 times

BiddlyBdoyng 3 years, 10 months ago

I think it's not this because Clouse Source Repositories is for source code. Sounds like we are looking for an executable?

upvoted 1 times

amxexam 4 years, 3 months ago

C&D we are all eliminating becoz of source storage repo Between A& B B looks more tempting to select because it mentions fire wallrule But the problem with B is the statement is wrong the access will happen from VM to storage and the statement mentions traffic from storage to Vm. Hence A

upvoted 3 times

Question #55 *Topic 1*

Your company is moving 75 TB of data into Google Cloud. You want to use Cloud Storage and follow Google-recommended practices. What should you do?

A. Move your data onto a Transfer Appliance. Use a Transfer Appliance Rehydrator to decrypt the data into Cloud Storage.

B. Move your data onto a Transfer Appliance. Use Cloud Dataprep to decrypt the data into Cloud Storage.

C. Install gsutil on each server that contains data. Use resumable transfers to upload the data into Cloud Storage.

D. Install gsutil on each server containing data. Use streaming transfers to upload the data into Cloud Storage.

Suggested Answer: *A*
Official Answer: A

Community vote distribution

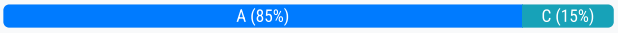

AshishK Highly Voted 6 years, 8 months ago
It should be 'A' Transfer Appliance lets you quickly and securely transfer large amounts of data to Google Cloud Platform via a high capacity storage server that you lease from Google and ship to our datacenter. Transfer Appliance is recommended for data that exceeds 20 TB or would take more than a week to upload.
upvoted 36 times

AshishK Highly Voted 6 years, 8 months ago
It should be 'A' Transfer Appliance lets you quickly and securely transfer large amounts of data to Google Cloud Platform via a high capacity storage server that you lease from Google and ship to our datacenter. Transfer Appliance is recommended for data that exceeds 20 TB or would take more than a week to upload.where did u get that 20 TB number - can help to share link?Here is the link: https://cloud.google.com/transfer-appliance/docs/2.2/overviewBut that link mentions a few hundred terabytes to 1 petabyte not 20TB or did I read that incorrectly?The request transfer appliance UI seems to suggest that it is not cost effective under 20TB of data.To this date Transfer Appliance supported locations are only United States Canada European Union Norway Switzerland. What if data reside in a location other than this? C is the most convenience for this scenario.There is NO mention of region - you dont assume anything NOT mentioned in the question therefore - Ans =AWhy assume a scenario no provided in the question. We need to choose the best case scenario based on available information instead of making assumptions. So A should be good.stupid answer
upvoted 36 times

MyPractice 6 years, 7 months ago
where did u get that 20 TB number - can help to share link?
upvoted 1 times

MyPractice 6 years, 7 months ago
where did u get that 20 TB number - can help to share link?Here is the link: https://cloud.google.com/transfer-appliance/docs/2.2/overviewBut that link mentions a few hundred terabytes to 1 petabyte not 20TB or did I read that incorrectly?The request transfer appliance UI seems to suggest that it is not cost effective under 20TB of data.
upvoted 1 times

onashwani 5 years, 8 months ago
Here is the link: https://cloud.google.com/transfer-appliance/docs/2.2/overviewBut that link mentions a few hundred terabytes to 1 petabyte not 20TB or did I read that incorrectly?The request transfer appliance UI seems to suggest that it is not cost effective under 20TB of data.
upvoted 4 times

onashwani 5 years, 8 months ago
Here is the link: https://cloud.google.com/transfer-appliance/docs/2.2/overview
upvoted 4 times

gcp_learner 4 years, 8 months ago
But that link mentions a few hundred terabytes to 1 petabyte not 20TB or did I read that incorrectly?The request transfer appliance UI seems to suggest that it is not cost effective under 20TB of data.
upvoted 2 times

gcp_learner 4 years, 8 months ago
But that link mentions a few hundred terabytes to 1 petabyte not 20TB or did I read that incorrectly?

upvoted 2 times

mindhoc 2 years, 9 months ago

The request transfer appliance UI seems to suggest that it is not cost effective under 20TB of data.

upvoted 2 times

Yahowmy 6 years, 1 month ago

To this date Transfer Appliance supported locations are only United States Canada European Union Norway Switzerland. What if data reside in a location other than this? C is the most convenience for this scenario.

upvoted 11 times

Yahowmy 6 years, 1 month ago

To this date Transfer Appliance supported locations are only United States Canada European Union Norway Switzerland. What if data reside in a location other than this? C is the most convenience for this scenario.There is NO mention of region - you dont assume anything NOT mentioned in the question therefore - Ans =AWhy assume a scenario no provided in the question. We need to choose the best case scenario based on available information instead of making assumptions. So A should be good.stupid answer

upvoted 11 times

NoCrapEva 2 years, 6 months ago

There is NO mention of region - you dont assume anything NOT mentioned in the question therefore - Ans =A

upvoted 1 times

Ramheadhunter 3 years, 12 months ago

Why assume a scenario no provided in the question. We need to choose the best case scenario based on available information instead of making assumptions. So A should be good.

upvoted 6 times

ccpmad 2 years, 2 months ago

stupid answer

upvoted 2 times

KouShikyou 6 years, 10 months ago

Why not A?A is okIt is AA, anything over 10TB goes via appliance.I have moved 120 TB using gsutil- cost effectively!takes longer thoughAnswer is A.

upvoted 30 times

KouShikyou 6 years, 10 months ago

Why not A?

upvoted 30 times

tartar 6 years ago

A is ok

upvoted 9 times

kumarp6 5 years, 9 months ago

It is A

upvoted 2 times

nitinz 5 years, 5 months ago

A, anything over 10TB goes via appliance.

upvoted 14 times

nitinz 5 years, 5 months ago

A, anything over 10TB goes via appliance.I have moved 120 TB using gsutil- cost effectively!takes longer though

upvoted 14 times

Begum 3 years, 10 months ago

I have moved 120 TB using gsutil- cost effectively!takes longer though

upvoted 1 times

Begum 3 years, 10 months ago

I have moved 120 TB using gsutil- cost effectively!

upvoted 1 times

zr79 3 years, 10 months ago

takes longer though

upvoted 1 times

Koushick 5 years, 3 months ago

Answer is A.

upvoted 4 times

halifax Most Recent 1 year, 5 months ago

what is a Rehydrator? a cloud version of WinRAR (compress, encrypt, decrypt, and decompress the data)

upvoted 1 times

alihabib 1 year, 6 months ago

A - Although I dont understand the Rehydration concept, but DataPrep is not for decryption. and gsutil is limited only for file size < 10GB ....

upvoted 1 times

plumbig11 1 year, 7 months ago

75tb Move your data onto a Transfer Appliance. Use a Transfer Appliance Rehydrator to decrypt the data into Cloud Storage.

upvoted 1 times

kip21 2 years, 7 months ago

A - Correct

upvoted 1 times

AdityaGupta 2 years, 10 months ago

Sending 75 TB of data on reliable 1.5 Gbps line will take about 6 days to complete data transfer online at the same time it will consume entire bandwidth. Hence use of Transfer appliace is required.

upvoted 1 times

MaheshKaswan 3 years, 2 months ago

Opion A is partially correct as you would not use a Transfer Appliance Rehydrator to decrypt the data. The Transfer Appliance itself is used to encrypt and decrypt the data. Option C is correct.

upvoted 2 times

RVivek 3 years, 6 months ago

gsutil is recommanded for data size less than a TB. That rules out C and D B says decrypt data using Dataprep not sure this is possible. https://cloud.google.com/solutions/migration-to-google-cloud-transferring-your-large-datasets#transfer-options

upvoted 3 times

n_nana 3 years, 7 months ago

It will be more precise with info about the bandwidth Google says: The two main criteria to consider with Transfer Appliance are cost and speed. With reasonable network connectivity (for example, 1 Gbps), transferring 100 TB of data online takes over 10 days to complete. If this rate is acceptable, an online transfer is likely a good solution for your needs. If you only have a 100 Mbps connection (or worse from a remote location), the same transfer takes over 100 days. At this point, it's worth considering an offline-transfer option such as Transfer Appliance. So even for such 100 TB google choose between transfer appliance or online transfer. not going with gsutil at all. it is clear gsutil is suitable for small to medium size (less than 1 TB) so with no more details, google recommendation is A

upvoted 3 times

omermahgoub 3 years, 8 months ago

The correct answer is A: Move your data onto a Transfer Appliance. Use a Transfer Appliance Rehydrator to decrypt the data into Cloud Storage.

To move large amounts of data into Google Cloud, it is recommended to use Transfer Appliance. Transfer Appliance is a physical storage device that you can use to transfer large amounts of data to Google Cloud quickly and securely. Once you have moved your data onto a Transfer Appliance, you can use a Transfer Appliance Rehydrator to decrypt the data and load it into Cloud Storage.

upvoted 3 times

omermahgoub 3 years, 8 months ago

The correct answer is A: Move your data onto a Transfer Appliance. Use a Transfer Appliance Rehydrator to decrypt the data into Cloud Storage. To move large amounts of data into Google Cloud, it is recommended to use Transfer Appliance. Transfer Appliance is a physical storage device that you can use to transfer large amounts of data to Google Cloud quickly and securely. Once you have moved your data onto a Transfer Appliance, you can use a Transfer Appliance Rehydrator to decrypt the data and load it into Cloud Storage.Option B: Using Cloud Dataprep to decrypt the data into Cloud Storage is not a valid option, as Cloud Dataprep is a data preparation tool that does not support data transfer or decryption. Option C: Using resumable transfers to upload the data into Cloud Storage is not a recommended option for moving large amounts of data, as resumable transfers are designed for smaller data sets and may not be efficient for transferring large amounts of data. Option D: Using streaming transfers to upload the data into Cloud Storage is not a recommended option for moving large amounts of data, as streaming transfers are designed for transferring real-time data streams and may not be efficient for transferring large amounts of data. Therefore, the correct answer is A: Move your data onto a Transfer Appliance. Use a Transfer Appliance Rehydrator to decrypt the data into Cloud Storage.

upvoted 3 times

omermahgoub 3 years, 8 months ago

Option B: Using Cloud Dataprep to decrypt the data into Cloud Storage is not a valid option, as Cloud Dataprep is a data preparation tool that does not support data transfer or decryption. Option C: Using resumable transfers to upload the data into Cloud Storage is not a recommended option for moving large amounts of data, as resumable transfers are designed for smaller data sets and may not be efficient for transferring large amounts of data. Option D: Using streaming transfers to upload the data into Cloud Storage is not a recommended option for moving large amounts of data, as

streaming transfers are designed for transferring real-time data streams and may not be efficient for transferring large amounts of data. Therefore, the correct answer is A: Move your data onto a Transfer Appliance. Use a Transfer Appliance Rehydrator to decrypt the data into Cloud Storage.

upvoted 6 times

megumin 3 years, 9 months ago

ok for A

upvoted 1 times

AzureDP900 3 years, 10 months ago

Option A Use a Transfer Appliance Rehydrator

upvoted 1 times

minmin2020 3 years, 10 months ago

A. Move your data onto a Transfer Appliance. Use a Transfer Appliance Rehydrator to decrypt the data into Cloud Storage.

upvoted 1 times

sgofficial 4 years ago

A is correct ... https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#transfer-options 1. gsutil is for lessthan <1 TB data with enough bandwidth, so C and D can be eliminated 2. option b can be eliminated since dataprep for decription is not correct 3. so only left over is a and its offline transfer, since the question did not give any time line when the transfer to be completed

upvoted 3 times

amxexam 4 years, 3 months ago

I would eliminate B & C as the question clearly methos google recomendations.About 10 TB or canal to we need to me transfer appliance. Let's not worry about what regions Ohk for now. Between A & B B is over kill as transfer applion allows decryption. Hence A

upvoted 1 times

meokey 4 years, 4 months ago

Is Transfer Appliance suitable for me? "Your data size is greater than or equal to 10TB." https://cloud.google.com/transfer-appliance/docs/4.0/overview#suitability

upvoted 3 times

Question #56 *Topic 1*

You have an application deployed on Google Kubernetes Engine using a Deployment named echo-deployment. The deployment is exposed using a Service called echo-service. You need to perform an update to the application with minimal downtime to the application. What should you do?

A. Use kubectl set image deployment/echo-deployment

B. Use the rolling update functionality of the Instance Group behind the Kubernetes cluster

C. Update the deployment yaml file with the new container image. Use kubectl delete deployment/echo-deployment and kubectl create ג€"f

D. Update the service yaml file which the new container image. Use kubectl delete service/echo-service and kubectl create ג€"f

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (93%)

ffk Highly Voted 6 years, 9 months ago
A is correct. B is funnyA is okYes A is correctOnly logical answer is A.This question asked "You need to perform an update to the application with minimal downtime to the application."
upvoted 47 times

ffk Highly Voted 6 years, 9 months ago
A is correct.

B is funny
upvoted 47 times

tartar 6 years ago
A is ok
upvoted 14 times

kumarp6 5 years, 9 months ago
Yes A is correct
upvoted 3 times

nitinz 5 years, 5 months ago
Only logical answer is A.
upvoted 1 times

DrLu 6 years, 8 months ago
This question asked "You need to perform an update to the application with minimal downtime to the application."
upvoted 3 times

jcmoranp 6 years, 9 months ago
Correct is A
upvoted 13 times

Ral17 Most Recent 4 months, 2 weeks ago
A is correct. The command kubectl set image deployment/echo-deployment updates the container image in the deployment. Kubernetes will handle the rolling update for you, ensuring that some pods are always running during the update process.
upvoted 1 times

askith 1 year, 4 months ago
A is ok.
upvoted 1 times

gaufchamp 1 year, 4 months ago
The correct answer is: A. Use kubectl set image deployment/echo-deployment Explanation: Rolling Updates: Kubernetes deployments support rolling updates by default. This means that when you update the deployment, Kubernetes will gradually replace the old pods with new ones, ensuring minimal downtime. Why A is Correct: The command kubectl set image deployment/echo-deployment updates the container image in the deployment. Kubernetes will handle the rolling update for you, ensuring that some pods are always running during the update process.
upvoted 2 times

sim7243 1 year, 9 months ago

option A

upvoted 1 times

isa_pr 2 years, 1 month ago

It's A. As per K8s documentation: To update the image of the application to version 2, use the set image subcommand, followed by the deployment name and the new image version: kubectl set image deployments/kubernetes-bootcamp kubernetes-bootcamp=docker.io/jocatalin/kubernetes-bootcamp:v2 The command notified the Deployment to use a different image for your app and initiated a rolling update." "

upvoted 6 times

kip21 2 years, 7 months ago

Option C - is the best option to perform an update to an application deployed on Google Kubernetes Engine with minimal downtime because it provides control over the update process, ensures high availability, and minimizes disruption. Rolling update functionality can also be used but requires more effort to implement. Option A and Option D may result in downtime if the new image is incompatible with the existing application.

upvoted 1 times

thewalker 2 years, 9 months ago

A https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps#updating_an_application

upvoted 4 times

AdityaGupta 2 years, 10 months ago

C. Update the deployment yaml file with the new container image. Use kubectl delete deployment/echo-deployment and kubectl create ɜ€"f I agreed with omermahgoub with his explanation.delete.. minimal downtime.. A is correct ^_^U

upvoted 1 times

AdityaGupta 2 years, 10 months ago

Selected Answer: C

C. Update the deployment yaml file with the new container image. Use kubectl delete deployment/echo-deployment and kubectl create ɜ€"f

I agreed with omermahgoub with his explanation.

upvoted 1 times

cacharritos 2 years, 7 months ago

delete.. minimal downtime.. A is correct ^_^U

upvoted 1 times

vamgcp 3 years, 6 months ago

To perform an update to the application with minimal downtime on Google Kubernetes Engine (GKE), you can use a rolling update strategy, which involves updating the application incrementally, one pod at a time, while ensuring that the updated pods are functioning properly before updating the next set. Here's the general process: kubectl set image deployment/echo-deployment echo=

upvoted 4 times

roaming_panda 3 years, 7 months ago

https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps says rolling updates and mentions same command . So 100 % A

upvoted 9 times

omermahgoub 3 years, 8 months ago

The correct answer is C: Update the deployment yaml file with the new container image. Use kubectl delete deployment/echo-deployment and kubectl create –f . To perform an update to an application deployed on Google Kubernetes Engine with minimal downtime, you can follow these steps: Update the deployment yaml file with the new container image. Use the kubectl delete deployment/echo-deployment command to delete the existing deployment. Use the kubectl create –f command to create a new deployment using the updated yaml file. This process, known as a rolling update, allows you to update your application with minimal downtime by replacing the old version of the application with the new version one pod at a time, while ensuring that there is always at least one pod available to serve traffic.using kubectl set image deployment/deployment will not allow you to perform an update to the application with minimal downtime, even if the deployment is exposed using a Service. This command will update the image of the containers in the deployment, but it will not perform a rolling update. A rolling update allows you to update your application with minimal downtime by replacing the old version of the application with the new version one pod at a time, while ensuring that there is always at least one pod available to serve traffic. Without a rolling update, all of the pods in the deployment will be replaced at the same time, which may result in downtime for the application.Option A: Using kubectl set image deployment/echo-deployment will update the image of the containers in the deployment, but it will not perform a rolling update and may result in downtime for the application. Option B: Using the rolling update functionality of the Instance Group behind the Kubernetes cluster is not a valid option, as the rolling update functionality is used to update the instances in the instance group, not the containers in a deployment. Option D: Up

upvoted 3 times

omermahgoub 3 years, 8 months ago

The correct answer is C: Update the deployment yaml file with the new container image. Use kubectl delete deployment/echo-deployment and kubectl create –f .

To perform an update to an application deployed on Google Kubernetes Engine with minimal downtime, you can follow these steps:

Update the deployment yaml file with the new container image. Use the kubectl delete deployment/echo-deployment command to delete the existing deployment. Use the kubectl create –f command to create a new deployment using the updated yaml file. This process, known as a rolling update, allows you to update your application with minimal downtime by replacing the old version of the application with the new version one pod at a time, while ensuring that there is always at least one pod available to serve traffic.

upvoted 3 times

omermahgoub 3 years, 8 months ago

using kubectl set image deployment/deployment will not allow you to perform an update to the application with minimal downtime, even if the deployment is exposed using a Service. This command will update the image of the containers in the deployment, but it will not perform a rolling update. A rolling update allows you to update your application with minimal downtime by replacing the old version of the application with the new version one pod at a time, while ensuring that there is always at least one pod available to serve traffic. Without a rolling update, all of the pods in the deployment will be replaced at the same time, which may result in downtime for the application.

upvoted 1 times

omermahgoub 3 years, 8 months ago

Option A: Using kubectl set image deployment/echo-deployment will update the image of the containers in the deployment, but it will not perform a rolling update and may result in downtime for the application.

Option B: Using the rolling update functionality of the Instance Group behind the Kubernetes cluster is not a valid option, as the rolling update functionality is used to update the instances in the instance group, not the containers in a deployment.

Option D: Updating the service yaml file with the new container image and using kubectl delete service/echo-service and kubectl create –f is not a valid option, as the service is not responsible for running the application containers and updating the service will not update the application.

upvoted 3 times

omermahgoub 3 years, 8 months ago

Option A: Using kubectl set image deployment/echo-deployment will update the image of the containers in the deployment, but it will not perform a rolling update and may result in downtime for the application. Option B: Using the rolling update functionality of the Instance Group behind the Kubernetes cluster is not a valid option, as the rolling update functionality is used to update the instances in the instance group, not the containers in a deployment. Option D: Updating the service yaml file with the new container image and using kubectl delete service/echo-service and kubectl create –f is not a valid option, as the service is not responsible for running the application containers and updating the service will not update the application.A is correct, update template spec image in deployment yml will trigger rollout deployAnswer is A, you are wrong as hell, if you delete deployment its obvious app will face downtime

upvoted 3 times

CkWongCk 3 years, 6 months ago

A is correct, update template spec image in deployment yml will trigger rollout deploy

upvoted 1 times

_kartik_raj 2 years, 10 months ago

Answer is A, you are wrong as hell, if you delete deployment its obvious app will face downtime

upvoted 3 times

jasenmornin 3 years, 8 months ago

I think A is correct: B. I don't understand the objective of this option. C and D. These are eliminated because they involve suffering a downtime when the resources are eliminated, so they are not fulfilling one of the requirements.

upvoted 3 times

markus_de 3 years, 8 months ago

Example from official Kubernetes docu (for NGINX): kubectl set image deployment/nginx-deployment nginx=nginx:1.16.1
https://kubernetes.io/docs/concepts/workloads/controllers/deployment/

upvoted 5 times

megumin 3 years, 9 months ago

ok for A

upvoted 2 times

AzureDP900 3 years, 10 months ago

A is right -- kubectl set image deployment/echo-deployment

upvoted 1 times

Question #57 *Topic 1*

Your company is using BigQuery as its enterprise data warehouse. Data is distributed over several Google Cloud projects. All queries on BigQuery need to be billed on a single project. You want to make sure that no query costs are incurred on the projects that contain the data. Users should be able to query the datasets, but not edit them.
How should you configure users' access roles?

A. Add all users to a group. Grant the group the role of BigQuery user on the billing project and BigQuery dataViewer on the projects that contain the data.

B. Add all users to a group. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery user on the projects that contain the data.

C. Add all users to a group. Grant the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data.

D. Add all users to a group. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery jobUser on the projects that contain the data.

Suggested Answer: *C*
Official Answer: C

Community vote distribution
C (100%)

RitwickKumar Highly Voted 3 years, 6 months ago
Both A & C are correct but using the principle of least privileges C is the most appropriate. BigQuery User: (roles/bigquery.user) When applied to a dataset, this role provides the ability to read the dataset's metadata and list tables in the dataset. When applied to a project, this role also provides the ability to run jobs, including queries, within the project. A principal with this role can enumerate their own jobs, cancel their own jobs, and enumerate datasets within a project. Additionally, allows the creation of new datasets within the project; the creator is granted the BigQuery Data Owner role(roles/bigquery.dataOwner) on these new datasets. Lowest-level resources where you can grant this role: Dataset BigQuery Job User: (roles/bigquery.jobUser) Provides permissions to run jobs, including queries, within the project. Lowest-level resources where you can grant this role: Project Source: https://cloud.google.com/bigquery/docs/access-control
upvoted 29 times

kimharsh Highly Voted 4 years ago
C is the correct Answer , A is wrong because bq User Permission will allow you to edit the dataset, which is something that we don't want in this scenario. B and D is wrong because "You want to make sure that no query costs are incurred on the projects that contain the data" so you don't want users to fire quires on the Project that contains the dataset , hence the "dataViewer" permission https://cloud.google.com/bigquery/docs/access-controlIt seems that User Permission doesn't allow to edit data, isn't it?
upvoted 21 times

kratosmat 2 years, 11 months ago
It seems that User Permission doesn't allow to edit data, isn't it?
upvoted 1 times

plumbig11 Most Recent 1 year, 7 months ago
BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data.
upvoted 1 times

Edgo97 1 year, 12 months ago
The link to refer here: https://cloud.google.com/bigquery/docs/access-control
upvoted 1 times

SidsA 2 years, 10 months ago
The "roles/bigquery.jobUser" role provides the permission to run jobs, including querying, exporting and copying data, and creating views and materialized views. This role does not provide permissions to create, update, or delete BigQuery resources, such as datasets, tables, and models. Users with this role can only interact with BigQuery through jobs. The "roles/bigquery.User" role, on the other hand, provides the permission to create, update, and delete BigQuery resources, as well as run jobs. This role includes all the permissions of the "roles/bigquery.jobUser" role, and in addition allows users to manage BigQuery resources, such as creating datasets, tables, and models, and modifying their schema and access controls.
upvoted 3 times

jlambdan 2 years, 11 months ago

A is wrong because https://cloud.google.com/bigquery/docs/access-control#bigquery.user C is correct because https://cloud.google.com/bigquery/docs/access-control#bigquery.jobUser

upvoted 1 times

jay9114 3 years, 1 month ago

Important statements from the prompt 1. All queries need to be billed to a single project - one project that queries data stored on other projects. Let's call this our billing project. a. jobUser is the best role to satisfy this need, because it provides permission to run jobs and queries within a project. 2. Other projects is where the data resides. These projects don't need much access besides the ability to be viewed (not edited). a. The dataViewer role provide permission to read all datasets in the project.

upvoted 7 times

omermahgoub 3 years, 1 month ago

The correct answer is A: Add all users to a group. Grant the group the role of BigQuery user on the billing project and BigQuery dataViewer on the projects that contain the data.

To make sure that no query costs are incurred on the projects that contain the data and allow users to query the datasets but not edit them, you should follow these steps:

Add all users to a group. Grant the group the role of BigQuery user on the billing project. This will allow the group to run queries on BigQuery and incur costs on the billing project. Grant the group the role of BigQuery dataViewer on the projects that contain the data. This will allow the group to view the datasets and run queries on them, but not edit them.

upvoted 1 times

omermahgoub 3 years, 1 month ago

The correct answer is A: Add all users to a group. Grant the group the role of BigQuery user on the billing project and BigQuery dataViewer on the projects that contain the data. To make sure that no query costs are incurred on the projects that contain the data and allow users to query the datasets but not edit them, you should follow these steps: Add all users to a group. Grant the group the role of BigQuery user on the billing project. This will allow the group to run queries on BigQuery and incur costs on the billing project. Grant the group the role of BigQuery dataViewer on the projects that contain the data. This will allow the group to view the datasets and run queries on them, but not edit them.BigQuery User (roles/bigquery.user) When applied to a dataset, this role provides the ability to read the dataset's metadata and list tables in the dataset. When applied to a project, this role also provides the ability to run jobs, including queries, within the project. A principal with this role can enumerate their own jobs, cancel their own jobs, and enumerate datasets within a project. Additionally, allows the creation of new datasets within the project; the creator is granted the BigQuery Data Owner role (roles/bigquery.dataOwner) on these new datasets. Bigquery.user has potential to create a dataset inside the project and creates becomes owner of the dataset. This is not the requirement stated in the question scenario. Answer is COption B: Granting the group the roles of BigQuery dataViewer on the billing project and BigQuery user on the projects that contain the data will not allow the group to incur costs on the billing project and will not meet the requirements of the scenario. Option C: Granting the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data will not allow the group to incur costs on the billing project and will not meet the requirements of the scenario. Option D: Granting the gro

upvoted 1 times

Diwz 1 year, 10 months ago

BigQuery User (roles/bigquery.user) When applied to a dataset, this role provides the ability to read the dataset's metadata and list tables in the dataset. When applied to a project, this role also provides the ability to run jobs, including queries, within the project. A principal with this role can enumerate their own jobs, cancel their own jobs, and enumerate datasets within a project. Additionally, allows the creation of new datasets within the project; the creator is granted the BigQuery Data Owner role (roles/bigquery.dataOwner) on these new datasets. Bigquery.user has potential to create a dataset inside the project and creates becomes owner of the dataset. This is not the requirement stated in the question scenario. Answer is C

upvoted 1 times

omermahgoub 3 years, 1 month ago

Option B: Granting the group the roles of BigQuery dataViewer on the billing project and BigQuery user on the projects that contain the data will not allow the group to incur costs on the billing project and will not meet the requirements of the scenario. Option C: Granting the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data will not allow the group to incur costs on the billing project and will not meet the requirements of the scenario. Option D: Granting the group the roles of BigQuery dataViewer on the billing project and BigQuery jobUser on the projects that contain the data will not allow the group to incur costs on the billing project and will not meet the requirements of the scenario.

upvoted 1 times

omermahgoub 3 years, 1 month ago

The BigQuery Job User role (roles/bigquery.jobUser) and the BigQuery User role (roles/bigquery.user) have similar permissions, but they differ in the scope of their permissions.

The BigQuery Job User role grants users the ability to create and modify query jobs, but it does not grant them the ability to run queries or incur costs on the project. This role is intended for users who need to create and manage query jobs, but who should not be able to run queries or incur costs.

The BigQuery User role grants users the ability to run queries and incur costs on the project, in addition to the ability to create and modify query jobs. This role is intended for users who need to run queries and incur costs on the project, as well as create and manage query jobs.

upvoted 3 times

omermahgoub 3 years, 1 month ago

The BigQuery Job User role (roles/bigquery.jobUser) and the BigQuery User role (roles/bigquery.user) have similar permissions, but they differ in the scope of their permissions. The BigQuery Job User role grants users the ability to create and modify query jobs, but it does not grant them the ability to run queries or incur costs on the project. This role is intended for users who need to create and manage query jobs, but who should not be able to run queries or incur costs. The BigQuery User role grants users the ability to run queries and incur costs on the project, in addition to the ability to create and modify query jobs. This role is intended for users who need to run queries and incur costs on the project, as well as create and manage query jobs.Here is a summary of the differences between the BigQuery Job User role and the BigQuery User role: BigQuery Job User role (roles/bigquery.jobUser): Can create and modify query jobs Cannot run queries or incur costs on the project BigQuery User role (roles/bigquery.user): Can create and modify query jobs Can run queries and incur costs on the project If you want to grant users the ability to create and modify query jobs, but not run queries or incur costs on the project, you should use the BigQuery Job User role. If you want to grant users the ability to run queries and incur costs on the project, in addition to the ability to create and modify query jobs, you should use the BigQuery User role.

upvoted 3 times

omermahgoub 3 years, 1 month ago

Here is a summary of the differences between the BigQuery Job User role and the BigQuery User role: BigQuery Job User role (roles/bigquery.jobUser): Can create and modify query jobs Cannot run queries or incur costs on the project BigQuery User role (roles/bigquery.user): Can create and modify query jobs Can run queries and incur costs on the project If you want to grant users the ability to create and modify query jobs, but not run queries or incur costs on the project, you should use the BigQuery Job User role. If you want to grant users the ability to run queries and incur costs on the project, in addition to the ability to create and modify query jobs, you should use the BigQuery User role.

upvoted 2 times

AzureDP900 3 years, 4 months ago

C is right Add all users to a group. Grant the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data.

upvoted 1 times

minmin2020 3 years, 4 months ago

C. Add all users to a group. Grant the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data.

upvoted 1 times

Vedjha 3 years, 4 months ago

D is the answer: Cloud BigQuery Roles Cloud BigQuery IAM Roles BigQuery Admin - bigquery.* BigQuery Data Owner - bigquery.datasets.*, bigquery.models.*, bigquery.routines.*, bigquery.tables.* (Does NOT have access to Jobs!) BigQuery Data Editor - bigquery.tables. (create/delete/export/get/getData/getIamPolicy/ list/update/updateData/updateTag), bigquery.models.*, bigquery.routines.*, bigquery.datasets. (create/get/getIamPolicy/updateTag) BigQuery Data Viewer - get/list bigquery.(datasets/models/routines/tables) BigQuery Job User - bigquery.jobs.create BigQuery User - BigQuery Data Viewer + get/list (jobs, capacityCommitments, reservations etc) To see data, you need either BigQuery User or BigQuery Data Viewer roles You CANNOT see data with BigQuery Job User roles BigQuery Data Owner or Data Viewer roles do NOT have access to jobs!

upvoted 1 times

kimharsh 4 years ago

C is the correct Answer , A is wrong because bq User Permission will allow you to edit the dataset, which is something that we don't want in this scenario. B and D is wrong because "You want to make sure that no query costs are incurred on the projects that contain the data" so you don't want users to fire quires on the Project that contains the dataset , hence the "dataViewer" permission https://cloud.google.com/bigquery/docs/access-control

upvoted 1 times

victory108 4 years, 1 month ago

C. Add all users to a group. Grant the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data.

upvoted 4 times

LoveT 4 years, 1 month ago

C looks to be the correct answer

upvoted 2 times

HenkH 4 years, 1 month ago

JobUser is the correct terminology for bq. Only read access to data sources is required.

upvoted 1 times

HenkH 4 years, 1 month ago

bq is using jobs - so "user" isn't specific enough, jobuser is.

upvoted 2 times

HenkH 4 years, 1 month ago

bq is using jobs - so "user" isn't specific enough, jobuser is.Hence C

upvoted 2 times

elenamatay 4 years, 1 month ago

Hence C

upvoted 1 times

Question #58 *Topic 1*

You have developed an application using Cloud ML Engine that recognizes famous paintings from uploaded images. You want to test the application and allow specific people to upload images for the next 24 hours. Not all users have a Google Account. How should you have users upload images?

A. Have users upload the images to Cloud Storage. Protect the bucket with a password that expires after 24 hours.

B. Have users upload the images to Cloud Storage using a signed URL that expires after 24 hours.

C. Create an App Engine web application where users can upload images. Configure App Engine to disable the application after 24 hours. Authenticate users via Cloud Identity.

D. Create an App Engine web application where users can upload images for the next 24 hours. Authenticate users via Cloud Identity.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (100%)

jcmoranp Highly Voted 6 years, 10 months ago
Correct answer is B
upvoted 44 times

jcmoranp Highly Voted 6 years, 10 months ago
Correct answer is BB is okSigned URL ... B is correctB signed URL
upvoted 44 times

tartar 6 years ago
B is ok
upvoted 8 times

kumarp6 5 years, 9 months ago
Signed URL ... B is correct
upvoted 3 times

nitinz 5 years, 5 months ago
B signed URL
upvoted 3 times

MyPractice 6 years, 7 months ago
Ans B "When should you use a signed URL? In some scenarios, you might not want to require your users to have a Google account in order to access Cloud Storage" "Signed URLs contain authentication information in their query string, allowing users without credentials to perform specific actions on a resource" https://cloud.google.com/storage/docs/access-control/signed-urls
upvoted 29 times

sim7243 Most Recent 1 year, 9 months ago
Correct answer is B
upvoted 1 times

gun123 2 years, 7 months ago
Correct answer is B
upvoted 1 times

red_panda 3 years ago
B is the answer
upvoted 1 times

fussili 3 years, 4 months ago
The correct answer is B. A is not a good choice because it is not possible to set an expiration time for a password protected Cloud Storage bucket. This means that if a user had the password, they would be able to upload images to the bucket even after the 24 hour period has expired. B is the correct answer because a signed URL can be generated to allow specific users to upload images to Cloud Storage without requiring them to have a Google Account. The URL can be set to expire after 24 hours, which ensures that users can only upload images during the allowed time period. C is not the best choice because it involves creating an App Engine web application, which is more complex than using Cloud Storage with a signed URL.

Additionally, App Engine instances cannot be turned off programmatically, so it would not be possible to disable the application after 24 hours. D option is similar to option C, but it involves creating an App Engine web application. This would add unnecessary complexity to the solution, and it would not provide any additional benefits compared to using Cloud Storage with a signed URL.

upvoted 6 times

omermahgoub 3 years, 8 months ago

The correct answer is B: Have users upload the images to Cloud Storage using a signed URL that expires after 24 hours.

To allow specific users to upload images to Cloud Storage for testing your Cloud ML Engine application, and to not require all users to have a Google Account, you should use signed URLs. A signed URL is a URL that allows access to a specific resource in Cloud Storage, and that is only valid for a specified period of time.

To create a signed URL that expires after 24 hours, you can use the gsutil signurl command. For example:

Copy code gsutil signurl -d 24h service-account.json gs://bucket-name/object-name This will generate a signed URL that allows users to upload an object to the specified bucket with the specified name, and that will only be valid for 24 hours.

upvoted 3 times

omermahgoub 3 years, 8 months ago

The correct answer is B: Have users upload the images to Cloud Storage using a signed URL that expires after 24 hours. To allow specific users to upload images to Cloud Storage for testing your Cloud ML Engine application, and to not require all users to have a Google Account, you should use signed URLs. A signed URL is a URL that allows access to a specific resource in Cloud Storage, and that is only valid for a specified period of time. To create a signed URL that expires after 24 hours, you can use the gsutil signurl command. For example: Copy code gsutil signurl -d 24h service-account.json gs://bucket-name/object-name This will generate a signed URL that allows users to upload an object to the specified bucket with the specified name, and that will only be valid for 24 hours.Option A: Protecting the bucket with a password that expires after 24 hours would not be a secure or scalable solution, as it would require you to distribute the password to all users and to update the password every 24 hours. Option C: Creating an App Engine web application where users can upload images, and configuring App Engine to disable the application after 24 hours, would not allow users to upload images after the application is disabled. Option D: Creating an App Engine web application where users can upload images for the next 24 hours and authenticating users via Cloud Identity would not allow users to upload images if they do not have a Google Account.

upvoted 3 times

omermahgoub 3 years, 8 months ago

Option A: Protecting the bucket with a password that expires after 24 hours would not be a secure or scalable solution, as it would require you to distribute the password to all users and to update the password every 24 hours. Option C: Creating an App Engine web application where users can upload images, and configuring App Engine to disable the application after 24 hours, would not allow users to upload images after the application is disabled. Option D: Creating an App Engine web application where users can upload images for the next 24 hours and authenticating users via Cloud Identity would not allow users to upload images if they do not have a Google Account.

upvoted 1 times

megumin 3 years, 9 months ago

B is ok

upvoted 1 times

AzureDP900 3 years, 10 months ago

B is right, Signed URL's will help in this scnerio.

upvoted 1 times

minmin2020 3 years, 10 months ago

B. Have users upload the images to Cloud Storage using a signed URL that expires after 24 hours.

upvoted 1 times

mv2000 4 years, 1 month ago

On 06/30/2022 Exam.

upvoted 2 times

mygcpjourney2712 4 years, 4 months ago

signed url

upvoted 1 times

vincy2202 4 years, 7 months ago

B is the correct answer

upvoted 2 times

haroldbenites 4 years, 8 months ago

Go for B.

upvoted 1 times

MaxNRG 4 years, 9 months ago

B – Have users upload the images to Cloud Storage via signed URL which expires after 24 hours. Signed URL is a preferable way to allow something with limited timeframe, doesn't require the account

upvoted 1 times

MamthaSJ 5 years, 1 month ago

Answer is B

upvoted 3 times

rishab86 5 years, 2 months ago

Answer B A signed URL is a URL that provides limited permission and time to make a request. Signed URLs contain authentication information in their query string, allowing users without credentials to perform specific actions on a resource. When you generate a signed URL, you specify a user or service account which must have sufficient permission to make the request that the signed URL will make. After you generate a signed URL, anyone who possesses it can use the signed URL to perform specified actions, such as reading an object, within a specified period of time. When should you use a signed URL? In some scenarios, you might not want to require your users to have a Google account in order to access Cloud Storage, but you still want to control access using your application-specific logic.

upvoted 4 times

Question #59 *Topic 1*

Your web application must comply with the requirements of the European Union's General Data Protection Regulation (GDPR). You are responsible for the technical architecture of your web application. What should you do?

A. Ensure that your web application only uses native features and services of Google Cloud Platform, because Google already has various certifications and provides ɹ€pass-onɹ€ compliance when you use native features.

B. Enable the relevant GDPR compliance setting within the GCPConsole for each of the services in use within your application.

C. Ensure that Cloud Security Scanner is part of your test planning strategy in order to pick up any compliance gaps.

D. Define a design for the security of data in your web application that meets GDPR requirements.

Suggested Answer: *D*

Reference: https://www.mobiloud.com/blog/gdpr-compliant-mobile-app/

Community vote distribution

D (100%)

AWS56 Highly Voted 5 years, 1 month ago

Agree D

upvoted 17 times

AshokC Highly Voted 4 years, 5 months ago

D - https://cloud.google.com/security/gdpr The GDPR lays out specific requirements for businesses and organizations who are established in Europe or who serve users in Europe. It: Regulates how businesses can collect, use, and store personal data Builds upon current documentation and reporting requirements to increase accountability Authorizes fines on businesses who fail to meet its requirements

upvoted 16 times

omermahgoub Most Recent 2 years, 1 month ago

The correct answer is option D: Define a design for the security of data in your web application that meets GDPR requirements.

The General Data Protection Regulation (GDPR) is a comprehensive data protection law that applies to any company that processes the personal data of individuals in the European Union (EU). As the technical architect of your web application, it is your responsibility to ensure that the application is compliant with GDPR requirements.

upvoted 4 times

omermahgoub 2 years, 1 month ago

The correct answer is option D: Define a design for the security of data in your web application that meets GDPR requirements. The General Data Protection Regulation (GDPR) is a comprehensive data protection law that applies to any company that processes the personal data of individuals in the European Union (EU). As the technical architect of your web application, it is your responsibility to ensure that the application is compliant with GDPR requirements.Option A: While it is true that Google has various certifications and provides pass-on compliance when you use native features, simply using native features and services of Google Cloud Platform is not sufficient to ensure compliance with GDPR. You still need to implement appropriate controls and safeguards to protect personal data and meet GDPR requirements. Option B: Enabling the relevant GDPR compliance setting within the GCP console for each of the services in use within your application may help ensure compliance with GDPR, but it is not sufficient on its own. You still need to implement appropriate controls and safeguards to protect personal data and meet GDPR requirements. Option C: Using Cloud Security Scanner as part of your test planning strategy can help identify potential security vulnerabilities and compliance gaps in your web application, but it is not sufficient on its own to ensure compliance with GDPR. You still need to implement appropriate controls and safeguards to protect personal data and meet GDPR requirements.

upvoted 4 times

omermahgoub 2 years, 1 month ago

Option A: While it is true that Google has various certifications and provides pass-on compliance when you use native features, simply using native features and services of Google Cloud Platform is not sufficient to ensure compliance with GDPR. You still need to implement appropriate controls and safeguards to protect personal data and meet GDPR requirements. Option B: Enabling the relevant GDPR compliance setting within the GCP console for each of the services in use within your application may help ensure compliance with GDPR, but it is not sufficient on its own. You still need to implement appropriate controls and safeguards to protect personal data and meet GDPR requirements. Option C: Using Cloud Security Scanner as part of your test planning strategy can help identify potential security vulnerabilities and compliance gaps in your web application, but it is

not sufficient on its own to ensure compliance with GDPR. You still need to implement appropriate controls and safeguards to protect personal data and meet GDPR requirements.

upvoted 10 times

megumin 2 years, 3 months ago

D is ok

upvoted 1 times

AzureDP900 2 years, 4 months ago

Define a design for the security of data in your web application that meets GDPR requirements. D is right

upvoted 1 times

vincy2202 3 years, 1 month ago

D is the correct answer

upvoted 1 times

haroldbenites 3 years, 2 months ago

Go for D

upvoted 1 times

joe2211 3 years, 2 months ago

vote D

upvoted 1 times

MaxNRG 3 years, 3 months ago

D – Define a design for the security of data in your web app that meets GDPR requirements.

upvoted 1 times

MikeB19 3 years, 5 months ago

A is wrong D is correct. The q refers is “Microsoft sql” not “MySQL”. App replication in MSsql is achieved with Availability Groups within MSsql https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/overview-of-always-on-availability-groups-sql-server?view=sql-server-ver15

upvoted 1 times

MamthaSJ 3 years, 7 months ago

Answer is D

upvoted 3 times

victory108 3 years, 9 months ago

D. Define a design for the security of data in your web application that meets GDPR requirements.

upvoted 2 times

un 3 years, 9 months ago

D is correct

upvoted 1 times

Ausias18 3 years, 10 months ago

Answer is D

upvoted 2 times

lynx256 3 years, 10 months ago

D is ok

upvoted 1 times

CloudGenious 4 years ago

you should design your app such that they meet GDPR req. As a customer google cloud, GDPR should be part of protection strategy .So the ans is D.

upvoted 3 times

bnlcnd 4 years ago

There is really no real answer in the 4 choices. D is so generic that it cannot go wrong :)

upvoted 4 times

Question #60 *Topic 1*

You need to set up Microsoft SQL Server on GCP. Management requires that there's no downtime in case of a data center outage in any of the zones within a
GCP region. What should you do?

A. Configure a Cloud SQL instance with high availability enabled.

B. Configure a Cloud Spanner instance with a regional instance configuration.

C. Set up SQL Server on Compute Engine, using Always On Availability Groups using Windows Failover Clustering. Place nodes in different subnets.

D. Set up SQL Server Always On Availability Groups using Windows Failover Clustering. Place nodes in different zones.

Suggested Answer: *D*
Official Answer: D

Community vote distribution
A (54%) D (45%)

learningpv Highly Voted 6 years, 6 months ago
A seems correct. "... high availability (HA) configuration for Cloud SQL instances... A Cloud SQL instance configured for HA is also called a regional instance and is located in a primary and secondary zone within the configured region. In the event of an instance or zone failure, this configuration reduces downtime, and your data continues to be available to client applications."
upvoted 66 times

learningpv Highly Voted 6 years, 6 months ago
A seems correct. "... high availability (HA) configuration for Cloud SQL instances... A Cloud SQL instance configured for HA is also called a regional instance and is located in a primary and secondary zone within the configured region. In the event of an instance or zone failure, this configuration reduces downtime, and your data continues to be available to client applications."It applies for MySQL and HA is not available for MS SQLYes it is available, its in beta, but when creating a "SQL Server 2017 Standard" in Cloud SQL menu you can chose single one or HA (regional).The problem is that these questions are ancient (talking about StackDriver all the time for example, it was rebranded in 2020!!!). So unfortunately we need to think of "What did the professor think 2-4 years ago" when this question was created. Otherwise I'd say A is the best all day!It is available, please see; https://cloud.google.com/sql/docs/sqlserver/high-availability?_ga=2.30855355.-503483612.1582800507 Also a video from Google; https://youtu.be/vMUpNoukwnMD is correct. Question is - "no downtime while installing MS SQL" , not on choosing or replacing with GCP product. I agree A is good solution for this requirements.. however question is not on choosing database.. its for HA.. so I will choose D.you said it yourself,"this configuration reduces downtime", it doesn't remove it completely as required, so the correct answer is D. See @taer answer: "...Cloud SQL's high availability configuration relies on a primary instance and a standby instance that resides in a different zone within the same region. In the event of an outage, Cloud SQL will automatically failover to the standby instance, but the failover process takes some time, usually around 1-2 minutes, which might result in a brief downtime. On the other hand, using SQL Server Always On Availability Groups with Windows Failover Clustering allows you to have multiple replicas of your data distributed across different zones. The failover pro
upvoted 66 times

learningpv 6 years, 6 months ago
It applies for MySQL and HA is not available for MS SQL
upvoted 5 times

learningpv 6 years, 6 months ago
It applies for MySQL and HA is not available for MS SQLYes it is available, its in beta, but when creating a "SQL Server 2017 Standard" in Cloud SQL menu you can chose single one or HA (regional).The problem is that these questions are ancient (talking about StackDriver all the time for example, it was rebranded in 2020!!!). So unfortunately we need to think of "What did the professor think 2-4 years ago" when this question was created. Otherwise I'd say A is the best all day!It is available, please see; https://cloud.google.com/sql/docs/sqlserver/high-availability?_ga=2.30855355.-503483612.1582800507 Also a video from Google; https://youtu.be/vMUpNoukwnM
upvoted 5 times

Jos 6 years, 6 months ago
Yes it is available, its in beta, but when creating a "SQL Server 2017 Standard" in Cloud SQL menu you can chose single one or HA (regional).
upvoted 3 times

Jos 6 years, 6 months ago

Yes it is available, its in beta, but when creating a "SQL Server 2017 Standard" in Cloud SQL menu you can chose single one or HA (regional).The problem is that these questions are ancient (talking about StackDriver all the time for example, it was rebranded in 2020!!!). So unfortunately we need to think of "What did the professor think 2-4 years ago" when this question was created. Otherwise I'd say A is the best all day!

upvoted 3 times

tocsa 2 years, 2 months ago

The problem is that these questions are ancient (talking about StackDriver all the time for example, it was rebranded in 2020!!!). So unfortunately we need to think of "What did the professor think 2-4 years ago" when this question was created. Otherwise I'd say A is the best all day!

upvoted 1 times

cetanx 6 years, 1 month ago

It is available, please see; https://cloud.google.com/sql/docs/sqlserver/high-availability?_ga=2.30855355.-503483612.1582800507 Also a video from Google; https://youtu.be/vMUpNoukwnM

upvoted 11 times

AmitAr 4 years, 3 months ago

D is correct. Question is - "no downtime while installing MS SQL" , not on choosing or replacing with GCP product. I agree A is good solution for this requirements.. however question is not on choosing database.. its for HA.. so I will choose D.

upvoted 4 times

DevOpsifier 3 years, 2 months ago

you said it yourself,"this configuration reduces downtime", it doesn't remove it completely as required, so the correct answer is D. See @taer answer: "...Cloud SQL's high availability configuration relies on a primary instance and a standby instance that resides in a different zone within the same region. In the event of an outage, Cloud SQL will automatically failover to the standby instance, but the failover process takes some time, usually around 1-2 minutes, which might result in a brief downtime.

On the other hand, using SQL Server Always On Availability Groups with Windows Failover Clustering allows you to have multiple replicas of your data distributed across different zones. The failover process between these replicas is faster and more seamless, reducing the risk of downtime even further..."

upvoted 6 times

DevOpsifier 3 years, 2 months ago

you said it yourself,"this configuration reduces downtime", it doesn't remove it completely as required, so the correct answer is D. See @taer answer: "...Cloud SQL's high availability configuration relies on a primary instance and a standby instance that resides in a different zone within the same region. In the event of an outage, Cloud SQL will automatically failover to the standby instance, but the failover process takes some time, usually around 1-2 minutes, which might result in a brief downtime. On the other hand, using SQL Server Always On Availability Groups with Windows Failover Clustering allows you to have multiple replicas of your data distributed across different zones. The failover process between these replicas is faster and more seamless, reducing the risk of downtime even further..."You can never have 100% uptime even Google advise against setting such standards/SLA. Going through GCP docs and youtube series, I have never come across "SQL Server Always On Availability Groups with Windows Failover Clustering", which makes me think the answer is A.

upvoted 6 times

piiizu 2 years, 10 months ago

You can never have 100% uptime even Google advise against setting such standards/SLA. Going through GCP docs and youtube series, I have never come across "SQL Server Always On Availability Groups with Windows Failover Clustering", which makes me think the answer is A.

upvoted 4 times

diluviouniv 5 years, 1 month ago

but it says: you need to setup SQL Server

upvoted 12 times

SMS 6 years, 5 months ago

Answer is A. Cloud SQL supports SQL Server and selecting high availability provides automatic failover within a region.

upvoted 32 times

hitmax87 Most Recent 2 months, 2 weeks ago

Configure a Cloud SQL instance with high availability enabled: While Cloud SQL for SQL Server offers an automated High Availability (HA) option, it does not guarantee "no downtime". During a zonal failover, Cloud SQL has to unbind the persistent disk from the primary VM, spin up the standby VM in the healthy zone, reattach the disk, and run database recovery. This process creates a noticeable failover delay/downtime window (typically taking 60 to 120 seconds), which violates management's explicit constraint.

upvoted 1 times

Quidde 3 months, 1 week ago

Gemini said D, and it was talking about MS SQL server.

upvoted 1 times

Ral17 4 months, 2 weeks ago

If the requirement is NO downtime, Cloud SQL does not meet the standard because of its failover delay. Instead, Always On Availability Groups on Compute Engine provide the most reliable, enterprise-grade SQL Server HA solution on GCP.

upvoted 1 times

Pime13 6 months ago

While Cloud SQL for SQL Server (Option A) is now a highly mature and popular option, it has historically been a "beta" or newer feature compared to the long-standing best practice of self-managing clusters for absolute control. Cloud SQL (Option A) is best for operational simplicity—Google handles patching, backups, and failover. It offers an SLA up to 99.99% for Enterprise Plus. Compute Engine (Option D) is best for full compatibility and specific SQL Server versions or features not yet available in the managed service.

upvoted 1 times

KCjoe 8 months ago

Managed High Availability: Cloud SQL for SQL Server provides a managed High Availability (HA) configuration, also known as a "regional instance". Regional Failover: In an HA setup, Cloud SQL maintains a primary instance in one zone and a standby instance in a different zone within the same region. Synchronous Replication: Data is synchronously replicated to persistent disks in both zones. If the primary zone experiences an outage, Cloud SQL automatically fails over to the standby instance in the healthy zone. Zero Application Changes: The application connects using the same static IP address, so no configuration changes are needed during a failover.

upvoted 2 times

JoseCloudEng1994 11 months, 3 weeks ago

Cloud SQL for SQL Server does have a built-in high availability feature that handles zonal outages. It automatically provisions a standby instance in a different zone. While this is a valid and often recommended solution for managed services, the question asks what you should do, and option D is a more direct and detailed answer to the technical setup using the specified technologies (SQL Server Always On). Both A and D are valid architectures for the problem, but D is a direct technical description of the solution.

upvoted 1 times

Jorgehr7 1 year, 1 month ago

Cloud SQL already allows deploy SQL Server

upvoted 2 times

Ankit267 1 year, 1 month ago

A as of today, A seems best, D is ok but A satisfies

upvoted 2 times

izekc 1 year, 5 months ago

As of now, A become most accurate answer already

upvoted 4 times

cd6aff3 1 year, 5 months ago

If your requirement is NO downtime, Cloud SQL does not meet the standard because of its failover delay. Instead, Always On Availability Groups (AGs) on Compute Engine provide the most reliable, enterprise-grade SQL Server HA solution on GCP.

upvoted 3 times

jobolesonihal 1 year, 5 months ago

While Cloud SQL for SQL Server does offer high availability options, it does not fully meet the requirements specified in the original query. Here's why: 1. Cloud SQL for SQL Server with high availability does provide protection against zonal outages, but it is limited to 99.99% availability for Enterprise Plus edition. 2. The original query specifically mentions "no downtime in case of a data center outage in any of the zones within a GCP region." Cloud SQL's high availability may still have some minimal downtime during failover. 3. The requirement for Microsoft SQL Server suggests a need for more control and customization than what Cloud SQL offers. The most appropriate solution that fully addresses the requirements is: Set up SQL Server Always On Availability Groups using Windows Failover Clustering. Place nodes in different zones (Option D).

upvoted 3 times

david_tay 1 year, 5 months ago

D is right. B is wrong as Microsoft SQL not supported by Spanner. A will have slight downtime when the standby instance is activated. Question clearly stated that no downtime is acceptable.

upvoted 3 times

Just_Ninja 1 year, 6 months ago

Cloud SQL Enterprise Plus-Version

upvoted 3 times

biplabid 1 year, 6 months ago

D is the correct answer. Although Cloud SQL supports SQL Server but not all features are available. One such feature is Always On Availability Group (AOAG) which ensures HA within a region.

upvoted 2 times

Question #61 *Topic 1*

The development team has provided you with a Kubernetes Deployment file. You have no infrastructure yet and need to deploy the application. What should you do?

A. Use gcloud to create a Kubernetes cluster. Use Deployment Manager to create the deployment.

B. Use gcloud to create a Kubernetes cluster. Use kubectl to create the deployment.

C. Use kubectl to create a Kubernetes cluster. Use Deployment Manager to create the deployment.

D. Use kubectl to create a Kubernetes cluster. Use kubectl to create the deployment.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (100%)

MeasService Highly Voted 6 years, 4 months ago
It has to be B. gcloud for creating cluster and kubectl for creating deployment
upvoted 54 times

KouShikyou Highly Voted 6 years, 4 months ago
May I ask why C is correct? I thought B was correct.
upvoted 26 times

KouShikyou 6 years, 4 months ago
May I ask why C is correct? I thought B was correct.B is correct, when you create a nodes in GKE you use gcloud rather than kubectl...B, gcloud to manage GKE and to manage pods use kubctl.yeap, gcloud command to create K8s cluster https://cloud.google.com/kubernetes-engine/docs/how-to/creating-a-clusterB is ok
upvoted 26 times

kumarp6 5 years, 3 months ago
B is correct, when you create a nodes in GKE you use gcloud rather than kubectl...
upvoted 4 times

nitinz 4 years, 11 months ago
B, gcloud to manage GKE and to manage pods use kubctl.
upvoted 2 times

res3 5 years, 7 months ago
yeap, gcloud command to create K8s cluster https://cloud.google.com/kubernetes-engine/docs/how-to/creating-a-clusterB is ok
upvoted 4 times

res3 5 years, 7 months ago
yeap, gcloud command to create K8s cluster https://cloud.google.com/kubernetes-engine/docs/how-to/creating-a-cluster
upvoted 4 times

tartar 5 years, 6 months ago
B is ok
upvoted 9 times

simiramis221 Most Recent 2 years, 2 months ago
Answer is B %100
upvoted 1 times

vamgcp 3 years ago
Create a Google Kubernetes Engine (GKE) cluster: You can use the Google Cloud Console or the gcloud command-line tool to create a GKE cluster, which will provide the underlying infrastructure for running your application. Deploy the application to the cluster: You can use the kubectl command-line tool to apply the Kubernetes Deployment file provided by the development team to the cluster.kubectl apply -f deployment.yaml
upvoted 4 times

megumin 3 years, 3 months ago
B is ok
upvoted 1 times

Mahmoud_E 3 years, 3 months ago

is the correct answer https://cloud.google.com/kubernetes-engine/docs/how-to/creating-a-cluster

upvoted 1 times

Mahmoud_E 3 years, 3 months ago

B is the correct answer cluster https://cloud.google.com/kubernetes-engine/docs/how-to/creating-a-cluster

upvoted 1 times

AzureDP900 3 years, 4 months ago

B is right

upvoted 1 times

SAMBIT 3 years, 11 months ago

Kubctle comes live only when cluster has been created in the cloud console using cloud command

upvoted 2 times

ghadxx 4 years ago

Deployment Manager is used to automate the process of provisioning infrastructure. Therefore, gcloud and Deployment Manager do the same thing. Meanwhile, kubectl is used to run commands against an already created cluster.

upvoted 9 times

haroldbenites 4 years, 2 months ago

Go for B. gcloud for create clusters. kubectl is used when the cluster already has been created. For example to create deployments. Kubectl has configured a config file where is specified the default cluster.

upvoted 2 times

vincy2202 4 years, 2 months ago

B is correct

upvoted 1 times

Zinhle 4 years, 3 months ago

Hi all may someone please share the link for the bank of questions because I cannot seem to locate them. thank you

upvoted 1 times

MaxNRG 4 years, 3 months ago

B – use gcloud to create cluster, use kubectl to create a deployment. https://cloud.google.com/kubernetes-engine/docs/how-to/creating-a-zonal-cluster In fact, kubectl run creates a deployment. https://cloud.google.com/kubernetes-engine/docs/tutorials/hello-app

upvoted 2 times

ale183 4 years, 5 months ago

Question for all , do we know if only new questions are part of the bank for new exam? Have any of the old questions appeared on new exam?

upvoted 3 times

xaliq 4 years, 5 months ago

B is corrent

upvoted 1 times

Raja101 4 years, 5 months ago

Why not A ?

upvoted 3 times

Question #62 *Topic 1*

You need to evaluate your team readiness for a new GCP project. You must perform the evaluation and create a skills gap plan which incorporates the business goal of cost optimization. Your team has deployed two GCP projects successfully to date. What should you do?

A. Allocate budget for team training. Set a deadline for the new GCP project.

B. Allocate budget for team training. Create a roadmap for your team to achieve Google Cloud certification based on job role.

C. Allocate budget to hire skilled external consultants. Set a deadline for the new GCP project.

D. Allocate budget to hire skilled external consultants. Create a roadmap for your team to achieve Google Cloud certification based on job role.

Suggested Answer: *A*

Official Answer: A

Community vote distribution

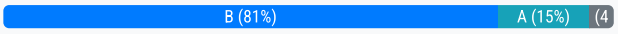

KouShikyou Highly Voted 6 years, 9 months ago

B is correct.

upvoted 46 times

KouShikyou Highly Voted 6 years, 9 months ago

B is correct.Yes it isI would agree with you because the question says create a skills gap planB, looks like cooked up question. Not gonna show up on actual test. Even if it does show up, its not market.exactly

upvoted 46 times

kumarp6 5 years, 9 months ago

Yes it is

upvoted 1 times

passnow 6 years, 8 months ago

I would agree with you because the question says create a skills gap plan

upvoted 4 times

nitinz 5 years, 5 months ago

B, looks like cooked up question. Not gonna show up on actual test. Even if it does show up, its not market.

upvoted 10 times

nitinz 5 years, 5 months ago

B, looks like cooked up question. Not gonna show up on actual test. Even if it does show up, its not market.exactly

upvoted 10 times

ACE_ASPIRE 4 years, 10 months ago

exactly

upvoted 1 times

mawsman 6 years, 6 months ago

I think it's B. "You must perform the evaluation and create a skills gap plan incorporates the business goal of cost optimization." The goal is to cost optimize - they might have deployed 2 projects but are they cost optimized? I think the only way to evaluate the skills gap in cost optimization is to make them get certified and use the results to determine cost optimization skills gap. Quickly pushing another project deadline would not help with cost optimization.Agreed. How is setting up a GCP project deadline helping towards skill gap and cost optimization.

upvoted 21 times

mawsman 6 years, 6 months ago

I think it's B. "You must perform the evaluation and create a skills gap plan incorporates the business goal of cost optimization." The goal is to cost optimize - they might have deployed 2 projects but are they cost optimized? I think the only way to evaluate the skills gap in cost optimization is to make them get certified and use the results to determine cost optimization skills gap. Quickly pushing another project deadline would not help with cost optimization.

upvoted 21 times

Smart 6 years, 6 months ago

Agreed. How is setting up a GCP project deadline helping towards skill gap and cost optimization.

upvoted 3 times

maxdanny Most Recent 1 year, 11 months ago

Explanation: Training and certification based on specific job roles (e.g., Cloud Architect, Data Engineer) will ensure your team has the necessary skills for the new GCP project and will help align their capabilities with cost optimization strategies. Since your team has already successfully deployed two GCP projects, upskilling them with targeted training is a more cost-effective and sustainable solution than relying on external consultants. This option balances both team readiness and the business goal of cost optimization, while building long-term internal expertise.

upvoted 4 times

joecloud12 2 years ago

b is obvious..u need to address the skill gap

upvoted 1 times

geekywitcher 2 years, 1 month ago

B is correct.

upvoted 1 times

Robert0 2 years, 2 months ago

B makes the most sense. If readiness is the goal, it makes sense to invest in formation

upvoted 1 times

GianpiGale 2 years, 3 months ago

C and D excluded because the teams have successfully deployed 2 GCP project to date, seems best option is to train them

upvoted 1 times

Teckexam 2 years, 7 months ago

B: this is logical answer. Setting project deadline or getting external consultants is not a good planning option.

upvoted 1 times

JPA210 2 years, 10 months ago

Answer A is more realistic. Certification is important, but not to be in a question in the exam.

upvoted 1 times

AdityaGupta 2 years, 10 months ago

You need to evaluate your team readiness for a new GCP project. You must perform the evaluation and create a skills gap plan which incorporates the business goal of cost optimization. The business Goal (Long Term Goal) is cost optimization, hence investment will be needed. Option A says setting up deadlines, it will not help in cost optimization. But giving training and obtaining certification will do.

upvoted 2 times

piiizu 2 years, 10 months ago

In the short run, A seems like the best option because this would enable a quicker transition to the cloud which is the business Goal of cost optimization. In terms of cost, certification will cost the organisation more. In the long run, B is the right option I would recommend as an architect. This would reduce the skill gap, increase proficiency and ensure repeatability. The organisation will incur more but will reap the reward... Except they have a bad culture and the staff resign right after getting the certification. Lol.

upvoted 2 times

Frusci 2 years, 11 months ago

B, better to train your team than hire, and setting a deadline would defeat the purpose of evaluating the team's readiness.

upvoted 2 times

PKookNN 3 years, 2 months ago

B is a better answer

upvoted 2 times

kaaiden 3 years, 3 months ago

GCP partnership need 3 project

upvoted 1 times

TheCloudGuruu 3 years, 3 months ago

B is the right answer

upvoted 2 times

grejao 3 years, 4 months ago

It shows more like a trick than a really question. Furthermore are this question relevant for certification?

upvoted 1 times

telp 3 years, 5 months ago

You are in a google exam. Always choose certification for your teams.

upvoted 3 times

Question #63 *Topic 1*

You are designing an application for use only during business hours. For the minimum viable product release, you'd like to use a managed product that automatically `scales to zero` so you don't incur costs when there is no activity.
Which primary compute resource should you choose?

A. Cloud Functions

B. Compute Engine

C. Google Kubernetes Engine

D. AppEngine flexible environment

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (96%) (4

abirroy Highly Voted 3 years, 11 months ago

Selected Answer: A

A. Cloud Functions - managed service scales down to 0 B. Compute Engine - not a managed service C. Google Kubernetes Engine - not a managed service and wont scale down to 0 D. AppEngine flexible environment - managed service but wont scale down to 0

upvoted 30 times

abirroy Highly Voted 3 years, 11 months ago

A. Cloud Functions - managed service scales down to 0 B. Compute Engine - not a managed service C. Google Kubernetes Engine - not a managed service and wont scale down to 0 D. AppEngine flexible environment - managed service but wont scale down to 0Agree with AGKE is a managed service.

upvoted 30 times

AzureDP900 3 years, 10 months ago

Agree with A

upvoted 3 times

NiveusSol 2 years, 5 months ago

GKE is a managed service.

upvoted 6 times

victory108 4 years, 7 months ago

A. Cloud Functions

upvoted 11 times

victory108 4 years, 7 months ago

A. Cloud FunctionsCloud function is more for event driven computing. We surely need k8s or app engine. Flex always have 1 instance running. So GKE should be the optionfrom the doc :https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler Note: If you specify a minimum of zero nodes, an idle node pool can scale down completely. However, at least one node must always be available in the cluster to run system Pods.But no cost for System/Control nodes

upvoted 11 times

vpatiltech 4 years, 6 months ago

Cloud function is more for event driven computing. We surely need k8s or app engine. Flex always have 1 instance running. So GKE should be the optionfrom the doc :https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler Note: If you specify a minimum of zero nodes, an idle node pool can scale down completely. However, at least one node must always be available in the cluster to run system Pods.But no cost for System/Control nodes

upvoted 6 times

vpatiltech 4 years, 6 months ago

Cloud function is more for event driven computing. We surely need k8s or app engine. Flex always have 1 instance running. So GKE should be the option

upvoted 6 times

YAS007 4 years, 3 months ago

from the doc :https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler

Note: If you specify a minimum of zero nodes, an idle node pool can scale down completely. However, at least one node must always be available in the cluster to run system Pods.

upvoted 1 times

YAS007 4 years, 3 months ago

from the doc :https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler Note: If you specify a minimum of zero nodes, an idle node pool can scale down completely. However, at least one node must always be available in the cluster to run system Pods.But no cost for System/Control nodes

upvoted 1 times

6721sora 3 years, 12 months ago

But no cost for System/Control nodes

upvoted 1 times

halifax Most Recent 1 year, 5 months ago

There is no correct answer, but I think, they meant to say cloud run (not cloud function). Cloud functions do not inherently scale to zero in the same way as Cloud Run.

upvoted 2 times

Toothpick 2 years ago

The only correct answer here is cloud run which isn't listed as an option. (A) can scale to zero, but is mean't for event driven workloads,not full sized applications, especially in the routing department (B) Needs manual stoppage, managed groups also keep one instance alive always (C) You're billed for the control plane even if the node pools are empty (D) Flex can't scale down to zero currently

upvoted 4 times

MiguelMiguel 2 years, 9 months ago

This question is confused. It's seems that the answer is cloud function because is the only platform that can scale 0 natively. If you want to scale to zero k8s, you have to create a solution based in scheduler, function and pub sub. Compute Engine is not a managed service, and app engine flex doesn't scale to 0. I would go to A but it's not clear.

upvoted 1 times

LaxmanTiwari 3 years, 3 months ago

agree A

upvoted 1 times

grejao 3 years, 4 months ago

Selected answer: D I choosed D, it appears that we do not. have a right answer here. A. Cloud Functions - its more for event driven computing, not for full application B. Compute Engine - not a managed service C. Google Kubernetes Engine - not a managed service and wont scale down to 0 D. AppEngine flexible environment - Only Standard App Engine can scale to 0.

upvoted 3 times

grejao 3 years, 4 months ago

Selected answer: D I choosed D, it appears that we do not. have a right answer here. A. Cloud Functions - its more for event driven computing, not for full application B. Compute Engine - not a managed service C. Google Kubernetes Engine - not a managed service and wont scale down to 0 D. AppEngine flexible environment - Only Standard App Engine can scale to 0.yes, but only Standard environment, not flexible environment

upvoted 3 times

Bedmed 3 years, 3 months ago

yes, but only Standard environment, not flexible environment

upvoted 4 times

GCPAnji 3 years, 4 months ago

For an application that is only used during business hours and needs to scale to zero during periods of inactivity to minimize costs, a good choice would be a Function-as-a-Service (FaaS) product like AWS Lambda or Google Cloud Functions.

upvoted 1 times

jlambdan 3 years, 5 months ago

Selected Answer: B

A a cloud function is not an application

B compute engine via MIG you can use an autoscaler with a schedule. https://cloud.google.com/compute/docs/autoscaler/scaling-schedules You then can go from 0 to more instance when required

C K8s is two complex for this. You can have an autoscaler for the cluster in order to get the node number to 0, but it require the node to have no pods

running. So you have to configure your deployments and all your workload to scale to 0 too. Other interference will be pod affinity, anti-affinity, disruption budget or unmanaged pod preventing pod eviction from node. But if the pod is not evicted, the node cannot be deleted. "autoscaler respects scheduling and eviction rules set on Pods. These restrictions can prevent a node from being deleted by the autoscaler. " => https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler#scheduling-and-disruption

D flexible cannot scale to 0

upvoted 1 times

jlambdan 3 years, 5 months ago

A a cloud function is not an application B compute engine via MIG you can use an autoscaler with a schedule. https://cloud.google.com/compute/docs/autoscaler/scaling-schedules You then can go from 0 to more instance when required C K8s is two complex for this. You can have an autoscaler for the cluster in order to get the node number to 0, but it require the node to have no pods running. So you have to configure your deployments and all your workload to scale to 0 too. Other interference will be pod affinity, anti-affinity, disruption budget or unmanaged pod preventing pod eviction from node. But if the pod is not evicted, the node cannot be deleted. "autoscaler respects scheduling and eviction rules set on Pods. These restrictions can prevent a node from being deleted by the autoscaler. " => https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler#scheduling-and-disruption D flexible cannot scale to 0App Engine is not an application on its own either, nor B or C options. B and C are not managed products; "managed" means that all of the infrastructure work (such as setting up an autoscheduler or autoscaler) will be managed by Google and not by the user. Using an IaaS solution for a managed use case is already contradictory. App Engine Flex will always have at least 1 instance running. While Cloud Functions is not an application itself, it's the only resource that can scale to 0. Plus, the requirement states that they will have an application running during business hours, but they never mentioned that the app will only be running in GCP or in Cloud Functions. They could be running their app anywhere else, and only call Cloud Functions when needed.That is not the concept of managed: https://cloud.google.com/blog/topics/developers-practitioners/serverless-vs-fully-managed-whats-difference

upvoted 1 times

jrisl1991 2 years, 10 months ago

App Engine is not an application on its own either, nor B or C options. B and C are not managed products; "managed" means that all of the infrastructure work (such as setting up an autoscheduler or autoscaler) will be managed by Google and not by the user. Using an IaaS solution for a managed use case is already contradictory. App Engine Flex will always have at least 1 instance running. While Cloud Functions is not an application itself, it's the only resource that can scale to 0. Plus, the requirement states that they will have an application running during business hours, but they never mentioned that the app will only be running in GCP or in Cloud Functions. They could be running their app anywhere else, and only call Cloud Functions when needed.That is not the concept of managed: https://cloud.google.com/blog/topics/developers-practitioners/serverless-vs-fully-managed-whats-difference

upvoted 1 times

jrisl1991 2 years, 10 months ago

App Engine is not an application on its own either, nor B or C options. B and C are not managed products; "managed" means that all of the infrastructure work (such as setting up an autoscheduler or autoscaler) will be managed by Google and not by the user. Using an IaaS solution for a managed use case is already contradictory.

App Engine Flex will always have at least 1 instance running.

While Cloud Functions is not an application itself, it's the only resource that can scale to 0. Plus, the requirement states that they will have an application running during business hours, but they never mentioned that the app will only be running in GCP or in Cloud Functions. They could be running their app anywhere else, and only call Cloud Functions when needed.

upvoted 1 times

theBestStudent 2 years, 8 months ago

That is not the concept of managed: https://cloud.google.com/blog/topics/developers-practitioners/serverless-vs-fully-managed-whats-difference

upvoted 1 times

CGS22 3 years, 5 months ago

A. Cloud Functions

upvoted 2 times

WFCheong 3 years, 7 months ago

Why not D? App Engines also can scale down to zero when there is no activity. https://cloud.google.com/appengine/docs/the-appengine-environments#:~:text=Intended%20to%20run%20for%20free,when%20there%20is%20no%20traffic. Intended to run for free or at very low cost, where you pay only for what you need and when you need it. For example, your application can scale to 0 instances when there is no traffic.There are 2 mode for App engines, standard and flexible. The standard environment can scale from zero instances up to thousands very quickly. In contrast, the flexible environment must have at least one instance running for each active version and can take longer to scale out in response to traffic.

upvoted 3 times

WFCheong 3 years, 7 months ago

Why not D? App Engines also can scale down to zero when there is no activity. https://cloud.google.com/appengine/docs/the-appengine-environments#:~:text=Intended%20to%20run%20for%20free,when%20there%20is%20no%20traffic. Intended to run for free or at very low cost, where you pay only for what you need and when you need it. For example, your application can scale to 0 instances when there is no traffic.

upvoted 3 times

CkWongCk 3 years, 6 months ago

There are 2 mode for App engines, standard and flexible. The standard environment can scale from zero instances up to thousands very quickly. In contrast, the flexible environment must have at least one instance running for each active version and can take longer to scale out in response to traffic.

upvoted 4 times

thamaster 3 years, 7 months ago

the only to scale to 0 is A

upvoted 2 times

omermahgoub 3 years, 8 months ago

The correct answer is A. Cloud Functions. Cloud Functions is a serverless compute service that lets you run code without provisioning or managing infrastructure. One of the key benefits of using Cloud Functions is that it automatically scales to meet the demands of your workload and automatically scales down to zero when there is no activity. This means that you only pay for the compute resources that you consume, which can help to reduce costs when your application is not in use. Additionally, Cloud Functions is easy to use and allows you to deploy your code with minimal effort, making it a good choice for a minimum viable product release.

upvoted 3 times

backhand 4 years ago

vote A this is easy one. key word: managed product, scales to zero scale to zero: app engine standard, cloud function

upvoted 2 times

vijbabu 4 years, 1 month ago

Answer is B

upvoted 1 times

Dhiraj03 4 years, 2 months ago

Cloud Functions can scale to zero when not in use

upvoted 3 times

Sskhan 4 years, 7 months ago

Answer is A, As Option D App engine flexible can have minimum 1 instance active.

upvoted 4 times

Question #64 *Topic 1*

You are creating an App Engine application that uses Cloud Datastore as its persistence layer. You need to retrieve several root entities for which you have the identifiers. You want to minimize the overhead in operations performed by Cloud Datastore. What should you do?

A. Create the Key object for each Entity and run a batch get operation

B. Create the Key object for each Entity and run multiple get operations, one operation for each entity

C. Use the identifiers to create a query filter and run a batch query operation

D. Use the identifiers to create a query filter and run multiple query operations, one operation for each entity

Suggested Answer: *A*

Official Answer: A

Community vote distribution



shashu07 Highly Voted 5 years, 8 months ago

Correct Answer: A Create the Key object for each Entity and run a batch get operation https://cloud.google.com/datastore/docs/best-practices Use batch operations for your reads, writes, and deletes instead of single operations. Batch operations are more efficient because they perform multiple operations with the same overhead as a single operation. Firestore in Datastore mode supports batch versions of the operations which allow it to operate on multiple objects in a single Datastore mode call. Such batch calls are faster than making separate calls for each individual entity because they incur the overhead for only one service call. If multiple entity groups are involved, the work for all the groups is performed in parallel on the server side.

upvoted 50 times

shashu07 Highly Voted 5 years, 8 months ago

Correct Answer: A Create the Key object for each Entity and run a batch get operation https://cloud.google.com/datastore/docs/best-practices Use batch operations for your reads, writes, and deletes instead of single operations. Batch operations are more efficient because they perform multiple operations with the same overhead as a single operation. Firestore in Datastore mode supports batch versions of the operations which allow it to operate on multiple objects in a single Datastore mode call. Such batch calls are faster than making separate calls for each individual entity because they incur the overhead for only one service call. If multiple entity groups are involved, the work for all the groups is performed in parallel on the server side.works fine .. A is right

upvoted 50 times

AzureDP900 3 years, 4 months ago

works fine .. A is right

upvoted 2 times

AWS56 6 years, 1 month ago

Agree A

upvoted 7 times

halifax Most Recent 1 year, 5 months ago

An old question, no more datastore. in 2025 you can use Firestore in datastore mode.

upvoted 2 times

de1001c 1 year, 8 months ago

Keep in mind that datastore is discontinued, Firestore being the recommended alternative.

upvoted 4 times

don_v 2 years, 1 month ago

According to "A. Create the Key object for each Entity and run a batch get operation" which is wrong as the key is already created for each entity whenever it's persisted. I believe the correct answer is C, -- to use a bulk query (a.k.a. "batch" in their terms). You need a query with a criteria anyway to find a resultset, and not just a fetch by "get" operation to load by surrogate keys.

upvoted 3 times

vamgcp 3 years ago

By using the "lookup by key" API of Cloud Datastore, you can minimize the overhead in operations performed by Cloud Datastore and optimize the performance of your App Engine application. from google.cloud import datastore client = datastore.Client() keys = [client.key('EntityKind', id) for id in entity_ids] entities = client.get_multi(keys)

upvoted 3 times

omermahgoub 3 years, 1 month ago

A. Create the Key object for each Entity and run a batch get operation

To minimize the overhead in operations performed by Cloud Datastore, you should use the batch get operation to retrieve multiple entities in a single API call. To do this, you should create a Key object for each entity that you want to retrieve, then pass the Key objects to the batch get operation. This will allow you to retrieve multiple entities in a single API call, reducing the number of operations performed by Cloud Datastore and improving the efficiency of your application.

upvoted 6 times

omermahgoub 3 years, 1 month ago

A. Create the Key object for each Entity and run a batch get operation To minimize the overhead in operations performed by Cloud Datastore, you should use the batch get operation to retrieve multiple entities in a single API call. To do this, you should create a Key object for each entity that you want to retrieve, then pass the Key objects to the batch get operation. This will allow you to retrieve multiple entities in a single API call, reducing the number of operations performed by Cloud Datastore and improving the efficiency of your application.Option B, running multiple get operations, one operation for each entity, would not be an efficient way to retrieve the entities because it would require multiple API calls to Cloud Datastore, which would increase the overhead and decrease the efficiency of the application. Option C, using the identifiers to create a query filter and running a batch query operation, would not be an efficient way to retrieve the entities because it would require performing a query operation, which is generally more expensive than a get operation. Option D, using the identifiers to create a query filter and running multiple query operations, one operation for each entity, would not be an efficient way to retrieve the entities because it would require performing multiple query operations, which are generally more expensive than get operations.

upvoted 6 times

omermahgoub 3 years, 1 month ago

Option B, running multiple get operations, one operation for each entity, would not be an efficient way to retrieve the entities because it would require multiple API calls to Cloud Datastore, which would increase the overhead and decrease the efficiency of the application. Option C, using the identifiers to create a query filter and running a batch query operation, would not be an efficient way to retrieve the entities because it would require performing a query operation, which is generally more expensive than a get operation. Option D, using the identifiers to create a query filter and running multiple query operations, one operation for each entity, would not be an efficient way to retrieve the entities because it would require performing multiple query operations, which are generally more expensive than get operations.

upvoted 6 times

megumin 3 years, 3 months ago

A is ok

upvoted 1 times

Mahmoud_E 3 years, 3 months ago

A is correct https://cloud.google.com/datastore/docs/best-practices#api_calls

upvoted 2 times

RitwickKumar 3 years, 6 months ago

https://cloud.google.com/datastore/docs/concepts/entities#datastore-datastore-batch-lookup-python

upvoted 3 times

haroldbenites 4 years, 2 months ago

go for A.

upvoted 2 times

vincy2202 4 years, 2 months ago

A is the right answer

upvoted 2 times

joe2211 4 years, 2 months ago

vote A

upvoted 2 times

MaxNRG 4 years, 3 months ago

A – create a key object for each entity, and run a batch get operations. See Batch Operations section here: https://cloud.google.com/datastore/docs/concepts/entities var keys = new Key[] { _keyFactory.CreateKey(1), _keyFactory.CreateKey(2) }; var tasks = _db.Lookup(keys[0], keys[1]); 1 and 2 are identifiers of the Key. Check Key / Identifier definition on the same link (top of that page) Such batch calls are faster than making separate calls for each individual entity because they incur the overhead for only one service call.

upvoted 1 times

victory108 4 years, 9 months ago

A. Create the Key object for each Entity and run a batch get operation

upvoted 1 times

un 4 years, 9 months ago

A is correct

upvoted 1 times

Ausias18 4 years, 10 months ago

Answer is A

upvoted 1 times

Question #65 *Topic 1*

You need to upload files from your on-premises environment to Cloud Storage. You want the files to be encrypted on Cloud Storage using customer-supplied encryption keys. What should you do?

A. Supply the encryption key in a .boto configuration file. Use gsutil to upload the files.

B. Supply the encryption key using gcloud config. Use gsutil to upload the files to that bucket.

C. Use gsutil to upload the files, and use the flag --encryption-key to supply the encryption key.

D. Use gsutil to create a bucket, and use the flag --encryption-key to supply the encryption key. Use gsutil to upload the files to that bucket.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (52%) C (42%)

KouShikyou Highly Voted 6 years, 10 months ago
In GCP document, key could be configured in .boto. I didn't find information show gsutil suppots flag "--encryption-key".

https://cloud.google.com/storage/docs/encryption/customer-supplied-keys
upvoted 48 times

KouShikyou Highly Voted 6 years, 10 months ago
In GCP document, key could be configured in .boto. I didn't find information show gsutil suppots flag "--encryption-key". https://cloud.google.com/storage/docs/encryption/customer-supplied-keysA is okA is correct.boto file with encryption key, but it will works for individual users, every user should update their own .boto with same key. Also while retrieving you should use the same key to decryption.The documentation is here: https://cloud.google.com/storage/docs/encryption/using-customer-supplied-keys#upload-encrypt Option C is correct. You can upload a file using customer-supplied encryption with the command: gcloud storage cp SOURCE_DATA gs://BUCKET_NAME/OBJECT_NAME --encryption-key=YOUR_ENCRYPTION_KEYOption C doesn't say "use gcloud storage cp", it says "use gsutil".
upvoted 48 times

tartar 6 years ago
A is ok
upvoted 16 times

nitinz 5 years, 5 months ago
A is correct
upvoted 4 times

kumarp6 5 years, 9 months ago
.boto file with encryption key, but it will works for individual users, every user should update their own .boto with same key. Also while retrieving you should use the same key to decryption.
upvoted 3 times

JaimeMS 2 years, 2 months ago
The documentation is here: https://cloud.google.com/storage/docs/encryption/using-customer-supplied-keys#upload-encrypt Option C is correct. You can upload a file using customer-supplied encryption with the command: gcloud storage cp SOURCE_DATA gs://BUCKET_NAME/OBJECT_NAME --encryption-key=YOUR_ENCRYPTION_KEYOption C doesn't say "use gcloud storage cp", it says "use gsutil".
upvoted 13 times

JaimeMS 2 years, 2 months ago
The documentation is here: https://cloud.google.com/storage/docs/encryption/using-customer-supplied-keys#upload-encrypt

Option C is correct. You can upload a file using customer-supplied encryption with the command: gcloud storage cp SOURCE_DATA gs://BUCKET_NAME/OBJECT_NAME --encryption-key=YOUR_ENCRYPTION_KEY
upvoted 13 times

ryaryarya 1 year, 7 months ago
Option C doesn't say "use gcloud storage cp", it says "use gsutil".
upvoted 9 times

Eroc 6 years, 9 months ago

I agree, A.(https://cloud.google.com/storage/docs/gsutil/addlhelp/UsingEncryptionKeys#generating-customer-supplied-encryption-keys)

upvoted 18 times

hitmax87 Most Recent 2 months, 2 weeks ago

This is deprecated question. They decommissioned gsutil

upvoted 1 times

cloud_rider 1 year, 5 months ago

A is Correct. Here is the documentation -https://cloud.google.com/storage/docs/boto-gsutil

upvoted 2 times

09bd94b 1 year, 6 months ago

A is the only correct answer. Those pointing to 'C' need to realize that the flag "--encryption-key" works with *gcloud storage cp* and NOT with gsutil

upvoted 2 times

hpf97 1 year, 6 months ago

https://cloud.google.com/storage/docs/gsutil Even it is recommended to use gcloud, this is the lonely answer which covers all the lifecycle of the process : create the storage, use a CEK, and upload files. I also agree that files have not to be encrypted individually, but the encryption has to be done at the storage level, so it will be transparent for a future consumer service. Answer B would be correct but it is not only a question of supplying the key, but also to update the bucket. So it is not enough

upvoted 1 times

09bd94b 1 year, 7 months ago

Only A is 100% correct

upvoted 2 times

Lrzo 1 year, 7 months ago

C is wrong because gsutil does not have a --encryption-key option

upvoted 1 times

klayytech 1 year, 7 months ago

The correct answer is C. Here's why: Customer-Supplied Encryption Keys (CSEK): When using CSEK, you provide the encryption key yourself. Google doesn't store your key on their servers. You're responsible for managing and protecting it. gsutil and --encryption-key: The gsutil command-line tool is the primary way to interact with Cloud Storage. To use CSEK with gsutil, you use the --encryption-key flag directly with the upload command. This flag takes the base64 encoded encryption key as its valu

upvoted 2 times

rrope 1 year, 7 months ago

Customer-Supplied Encryption Keys (CSEK) are provided on a per-request basis. This means you provide the key during the upload operation itself, not when creating the bucket or through persistent configuration files. gsutil is the command-line tool for interacting with Cloud Storage. The --encryption-key flag specifically allows you to provide the base64 encoded encryption key when uploading objects.

upvoted 2 times

rahuld19 1 year, 7 months ago

right answer is A

upvoted 2 times

mahi_h 1 year, 8 months ago

I see option D is not even discussed. The question said "upload files", meaning multiple object. Isn't the encrypted bucked creation a secured way to store them in cloud storage?

upvoted 2 times

kip21 1 year, 8 months ago

[GSUtil] check_hashes content_language decryption_key1 ... 100 default_api_version disable_analytics_prompt encryption_key

upvoted 2 times

deep316 1 year, 8 months ago

Option C: Use gsutil to upload the files and use the flag --encryption-key to supply the encryption key. This is the correct approach, as it allows you to specify the CSEK directly at the time of upload, ensuring that your files are encrypted using your provided key.

upvoted 2 times

klayytech 1 year, 8 months ago

D. Use gsutil to create a bucket, and use the flag --encryption-key to supply the encryption key. Use gsutil to upload the files to that bucket. This option provides the most comprehensive and secure approach: Create an encrypted bucket: Use gsutil mb -b location gs://your-bucket-name --encryption-key=your_encryption_key This ensures that all objects uploaded to this bucket will be encrypted with your provided key. Upload files to the encrypted bucket: Use gsutil cp your_local_file gs://your-bucket-name By following this approach, you guarantee that your files are encrypted both at rest and in

transit on Cloud Storage, providing a robust security posture. The other options either lack the encryption key specification or do not create an encrypted bucket, leaving your data vulnerable.

upvoted 4 times

desertlotus1211 1 year, 8 months ago

Selected Answer: A

The boto configuration file in Google Cloud Platform (GCP) controls how the gsutil command behaves:

Setting up gsutil You can use the boto configuration file to set up gsutil to work through a proxy.

Using encryption keys You can use the boto configuration file to use customer-managed or customer-supplied encryption keys.

upvoted 2 times

desertlotus1211 1 year, 8 months ago

The boto configuration file in Google Cloud Platform (GCP) controls how the gsutil command behaves: Setting up gsutil You can use the boto configuration file to set up gsutil to work through a proxy. Using encryption keys You can use the boto configuration file to use customer-managed or customer-supplied encryption keys..boto is smoother to use consistently...

upvoted 2 times

desertlotus1211 1 year, 8 months ago

.boto is smoother to use consistently...

upvoted 1 times

icarogsm 1 year, 8 months ago

A! I agree that the boto file sounds better

upvoted 2 times

Question #66 *Topic 1*

Your customer wants to capture multiple GBs of aggregate real-time key performance indicators (KPIs) from their game servers running on Google Cloud Platform and monitor the KPIs with low latency. How should they capture the KPIs?

A. Store time-series data from the game servers in Google Bigtable, and view it using Google Data Studio.

B. Output custom metrics to Stackdriver from the game servers, and create a Dashboard in Stackdriver Monitoring Console to view them.

C. Schedule BigQuery load jobs to ingest analytics files uploaded to Cloud Storage every ten minutes, and visualize the results in Google Data Studio.

D. Insert the KPIs into Cloud Datastore entities, and run ad hoc analysis and visualizations of them in Cloud Datalab.

Suggested Answer: *B*

Reference: https://cloud.google.com/solutions/data-lifecycle-cloud-platform

Community vote distribution

B (94%)

suryalsp Highly Voted 6 years, 8 months ago

Ans is B. Data studio cannot be used with BigTable https://datastudio.google.com/datahttps://datastudio.google.com/dataA is correctAs of today you canSource? https://lookerstudio.google.com/data?search=bigThat's BigQuery, not BigTable, no?Looker is not the same as Data Studio. I know it came to replace it, but these questions are kind of old, so unless it clearly says "looker", I wouldn't take both as the same.

upvoted 34 times

suryalsp Highly Voted 6 years, 8 months ago

Ans is B. Data studio cannot be used with BigTable https://datastudio.google.com/datahttps://datastudio.google.com/data

upvoted 34 times

Raja101 4 years, 11 months ago

A is correct

upvoted 4 times

ErenYeager 3 years, 9 months ago

As of today you canSource? https://lookerstudio.google.com/data?search=bigThat's BigQuery, not BigTable, no?Looker is not the same as Data Studio. I know it came to replace it, but these questions are kind of old, so unless it clearly says "looker", I wouldn't take both as the same.

upvoted 7 times

ErenYeager 3 years, 9 months ago

As of today you can

upvoted 7 times

anshumankmr80 3 years, 8 months ago

Source? https://lookerstudio.google.com/data?search=bigThat's BigQuery, not BigTable, no?Looker is not the same as Data Studio. I know it came to replace it, but these questions are kind of old, so unless it clearly says "looker", I wouldn't take both as the same.

upvoted 2 times

anshumankmr80 3 years, 8 months ago

Source?

https://lookerstudio.google.com/data?search=big

upvoted 2 times

HD2023 3 years, 4 months ago

That's BigQuery, not BigTable, no?

upvoted 2 times

jrisl1991 2 years, 9 months ago

Looker is not the same as Data Studio. I know it came to replace it, but these questions are kind of old, so unless it clearly says "looker", I wouldn't take both as the same.

upvoted 2 times

kolcsarzs 6 years, 8 months ago

correct is B

upvoted 12 times

nareshthumma 1 year, 9 months ago

Answer is B

upvoted 1 times

belouh 1 year, 10 months ago

Bigtable is not real time solution

upvoted 1 times

kip21 2 years, 7 months ago

B - Correct A - It is not a real-time solution

upvoted 1 times

kip21 2 years, 7 months ago

B - Correct A - It is not a real-time solutionBigTable is real time solution, BigQuery is notWhy is A not a real time solution? https://cloud.google.com/bigtable/docs/integrations#opentsdb it says: OpenTSDB is a time-series database that can use Bigtable for storage. Monitoring time-series data with OpenTSDB on Bigtable and GKE shows how to use OpenTSDB to collect, record, and monitor time-series data on Google Cloud. The OpenTSDB documentation provides additional information to help you get started.

upvoted 1 times

hitmax87 Most Recent 2 months, 2 weeks ago

BigTable is real time solution, BigQuery is not

upvoted 1 times

Saxena_Vibhor 2 years, 6 months ago

Why is A not a real time solution? https://cloud.google.com/bigtable/docs/integrations#opentsdb it says: OpenTSDB is a time-series database that can use Bigtable for storage. Monitoring time-series data with OpenTSDB on Bigtable and GKE shows how to use OpenTSDB to collect, record, and monitor time-series data on Google Cloud. The OpenTSDB documentation provides additional information to help you get started.

upvoted 3 times

AdityaGupta 2 years, 10 months ago

BigTable doesn't integrate with Data Studio https://cloud.google.com/bigtable/docs/integrations

upvoted 1 times

LaxmanTiwari 3 years, 3 months ago

B is correct, you should create custom KPI in Stack Driver

upvoted 1 times

omermahgoub 3 years, 8 months ago

B. Output custom metrics to Stackdriver from the game servers, and create a Dashboard in Stackdriver Monitoring Console to view them. To capture multiple GBs of aggregate real-time KPIs from game servers running on Google Cloud Platform and monitor them with low latency, the customer should output custom metrics to Stackdriver from the game servers. Stackdriver allows you to collect and store custom metrics, as well as view and analyze them in real-time using the Stackdriver Monitoring Console. The customer can create a Dashboard in the Monitoring Console to view the KPIs and monitor them with low latency.Option A, storing time-series data in Bigtable and viewing it using Data Studio, would not be suitable for capturing and monitoring real-time KPIs with low latency. Bigtable is a scalable NoSQL database that is optimized for large-scale batch processing, and Data Studio is a visualization tool that is not designed for real-time data analysis. Option C, scheduling BigQuery load jobs to ingest analytics files uploaded to Cloud Storage every ten minutes and visualizing the results in Data Studio, would not be suitable for capturing and monitoring real-time KPIs with low latency. BigQuery is a data warehouse that is optimized for batch processing, and it is not designed for real-time data analysis. Option D, inserting the KPIs into Cloud Datastore entities and running ad hoc analysis and visualizations of them in Cloud Datalab, would not be suitable for capturing and monitoring real-time KPIs with low latency. Cloud Datastore is a NoSQL document database, and Cloud Datalab is a data analysis and visualization tool that is not designed for real-time data analysis.big table is not for batch. It is used in IOT... https://cloud.google.com/bigtable

upvoted 10 times

omermahgoub 3 years, 8 months ago

B. Output custom metrics to Stackdriver from the game servers, and create a Dashboard in Stackdriver Monitoring Console to view them.

To capture multiple GBs of aggregate real-time KPIs from game servers running on Google Cloud Platform and monitor them with low latency, the customer should output custom metrics to Stackdriver from the game servers. Stackdriver allows you to collect and store custom metrics, as well as view and analyze them in real-time using the Stackdriver Monitoring Console. The customer can create a Dashboard in the Monitoring Console to view the KPIs and monitor them with low latency.

upvoted 10 times

omermahgoub 3 years, 8 months ago

Option A, storing time-series data in Bigtable and viewing it using Data Studio, would not be suitable for capturing and monitoring real-time KPIs with low latency. Bigtable is a scalable NoSQL database that is optimized for large-scale batch processing, and Data Studio is a visualization tool that is

not designed for real-time data analysis. Option C, scheduling BigQuery load jobs to ingest analytics files uploaded to Cloud Storage every ten minutes and visualizing the results in Data Studio, would not be suitable for capturing and monitoring real-time KPIs with low latency. BigQuery is a data warehouse that is optimized for batch processing, and it is not designed for real-time data analysis. Option D, inserting the KPIs into Cloud Datastore entities and running ad hoc analysis and visualizations of them in Cloud Datalab, would not be suitable for capturing and monitoring real-time KPIs with low latency. Cloud Datastore is a NoSQL document database, and Cloud Datalab is a data analysis and visualization tool that is not designed for real-time data analysis.big table is not for batch. It is used in IOT... https://cloud.google.com/bigtable

upvoted 9 times

omermahgoub 3 years, 8 months ago

Option A, storing time-series data in Bigtable and viewing it using Data Studio, would not be suitable for capturing and monitoring real-time KPIs with low latency. Bigtable is a scalable NoSQL database that is optimized for large-scale batch processing, and Data Studio is a visualization tool that is not designed for real-time data analysis.

Option C, scheduling BigQuery load jobs to ingest analytics files uploaded to Cloud Storage every ten minutes and visualizing the results in Data Studio, would not be suitable for capturing and monitoring real-time KPIs with low latency. BigQuery is a data warehouse that is optimized for batch processing, and it is not designed for real-time data analysis.

Option D, inserting the KPIs into Cloud Datastore entities and running ad hoc analysis and visualizations of them in Cloud Datalab, would not be suitable for capturing and monitoring real-time KPIs with low latency. Cloud Datastore is a NoSQL document database, and Cloud Datalab is a data analysis and visualization tool that is not designed for real-time data analysis.

upvoted 9 times

jlambdan 3 years, 5 months ago

big table is not for batch. It is used in IOT... https://cloud.google.com/bigtable

upvoted 5 times

megumin 3 years, 9 months ago

B is ok

upvoted 1 times

Mahmoud_E 3 years, 10 months ago

B is correct as Data studio does not support bigtable as a source

upvoted 4 times

zr79 3 years, 10 months ago

KPI, SLO,SLI all those work with observability which stackdriver

upvoted 2 times

jay9114 3 years, 10 months ago

The reference provided seems irrelevant to this question.

upvoted 1 times

bolington 4 years, 2 months ago

A is the correct answer, the key word here, real time and low latency.

upvoted 1 times

Nirca 4 years, 3 months ago

BigTable has no connection to data studio. https://datastudio.google.com/data?search=Big

upvoted 6 times

BeetleJuice 4 years, 6 months ago

B, it is

upvoted 2 times

OrangeTiger 4 years, 7 months ago

I don't think there is a correct answer, but B looks correct in this. If use Bigquery, then A is correct . C is not for realtime. D Datastore is for small usecase. Keywords 'real time' ,'analytics' https://events.withgoogle.com/solution-design-pattern-gaming/analytics-pattern/

upvoted 2 times

Question #67 *Topic 1*

You have a Python web application with many dependencies that requires 0.1 CPU cores and 128 MB of memory to operate in production. You want to monitor and maximize machine utilization. You also want to reliably deploy new versions of the application. Which set of steps should you take?

A. Perform the following: 1. Create a managed instance group with f1-micro type machines. 2. Use a startup script to clone the repository, check out the production branch, install the dependencies, and start the Python app. 3. Restart the instances to automatically deploy new production releases.

B. Perform the following: 1. Create a managed instance group with n1-standard-1 type machines. 2. Build a Compute Engine image from the production branch that contains all of the dependencies and automatically starts the Python app. 3. Rebuild the Compute Engine image, and update the instance template to deploy new production releases.

C. Perform the following: 1. Create a Google Kubernetes Engine (GKE) cluster with n1-standard-1 type machines. 2. Build a Docker image from the production branch with all of the dependencies, and tag it with the version number. 3. Create a Kubernetes Deployment with the imagePullPolicy set to 'IfNotPresent' in the staging namespace, and then promote it to the production namespace after testing.

D. Perform the following: 1. Create a GKE cluster with n1-standard-4 type machines. 2. Build a Docker image from the master branch with all of the dependencies, and tag it with 'latest'. 3. Create a Kubernetes Deployment in the default namespace with the imagePullPolicy set to 'Always'. Restart the pods to automatically deploy new production releases.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
C (57%) A (34%) B (9%)

jcmoranp Highly Voted 6 years, 9 months ago
C is correct, need "ifnotpresent"when uploads to container registry
upvoted 41 times

jcmoranp Highly Voted 6 years, 9 months ago
C is correct, need "ifnotpresent"when uploads to container registryifnotpresent won't pull new version.yes i agree
upvoted 41 times

medi01 3 years, 4 months ago
ifnotpresent won't pull new version.
upvoted 4 times

heretolearnazure 2 years, 12 months ago
yes i agree
upvoted 1 times

TosO 6 years, 9 months ago
C is the best choice. You can create a k8s cluster with just one node and use a different namespaces for staging and production. In staging, you will test the changes
upvoted 24 times

TosO 6 years, 9 months ago
C is the best choice. You can create a k8s cluster with just one node and use a different namespaces for staging and production. In staging, you will test the changesAgreed
upvoted 24 times

AzureDP900 3 years, 10 months ago
Agreed
upvoted 1 times

JoseCloudEng1994 11 months, 3 weeks ago
Selected Answer: C
The core of the question is not just to get the application running, but to "monitor and maximize machine utilization" and to "reliably deploy new versions."

Maximizing Utilization: An application requiring just 0.1 CPU cores and 128 MB of memory is a perfect fit for a container-orchestration platform. A single n1-standard-1 machine has 1 vCPU and 3.75 GB of memory. In a Managed Instance Group (MIG) with option B, a single application uses only a fraction of those resources, leaving the rest idle. In a GKE cluster (Option C), Kubernetes' scheduler can "bin-pack" multiple pods from the same or different applications onto that same node, significantly increasing resource utilization and cost-efficiency.

upvoted 2 times

JoseCloudEng1994 11 months, 3 weeks ago

The core of the question is not just to get the application running, but to "monitor and maximize machine utilization" and to "reliably deploy new versions." Maximizing Utilization: An application requiring just 0.1 CPU cores and 128 MB of memory is a perfect fit for a container-orchestration platform. A single n1-standard-1 machine has 1 vCPU and 3.75 GB of memory. In a Managed Instance Group (MIG) with option B, a single application uses only a fraction of those resources, leaving the rest idle. In a GKE cluster (Option C), Kubernetes' scheduler can "bin-pack" multiple pods from the same or different applications onto that same node, significantly increasing resource utilization and cost-efficiency.Reliable Deployments: GKE provides a superior deployment pipeline. Using Kubernetes Deployments offers features like rolling updates, automated health checks, and rollbacks, which are more robust than the instance template updates in a MIG. The container-based workflow is also faster and more flexible for frequent updates. Why B and D Are Flawed Option B (MIG with VM Images): While a common approach, building and deploying new VM images for every minor change is a slower and less agile process compared to building and pushing a new container image. It's a less efficient deployment model. Option D (GKE with Bad Practices): This option is the worst of the bunch. Using an n1-standard-4 machine is massive over-provisioning, completely failing to meet the goal of maximizing utilization. The use of the mutable latest tag and imagePullPolicy: Always is a major anti-pattern in production, leading to unpredictable and irreproducible deployments.1. Inefficient Resource Utilization (f1-micro): The f1-micro machine type is a shared-core VM, meaning its CPU is a fraction of a physical core and offers burstable performance. The application needs a minimum of 0.1 CPU cores. While an f1-micro can provide this, it is not designed for sustained production workloads and can be unreliable

upvoted 2 times

JoseCloudEng1994 11 months, 3 weeks ago

Reliable Deployments: GKE provides a superior deployment pipeline. Using Kubernetes Deployments offers features like rolling updates, automated health checks, and rollbacks, which are more robust than the instance template updates in a MIG. The container-based workflow is also faster and more flexible for frequent updates. Why B and D Are Flawed Option B (MIG with VM Images): While a common approach, building and deploying new VM images for every minor change is a slower and less agile process compared to building and pushing a new container image. It's a less efficient deployment model. Option D (GKE with Bad Practices): This option is the worst of the bunch. Using an n1-standard-4 machine is massive over-provisioning, completely failing to meet the goal of maximizing utilization. The use of the mutable latest tag and imagePullPolicy: Always is a major anti-pattern in production, leading to unpredictable and irreproducible deployments.1. Inefficient Resource Utilization (f1-micro): The f1-micro machine type is a shared-core VM, meaning its CPU is a fraction of a physical core and offers burstable performance. The application needs a minimum of 0.1 CPU cores. While an f1-micro can provide this, it is not designed for sustained production workloads and can be unreliable under load. This machine type is best suited for low-traffic websites or small-scale testing, not a production application that needs to maximize utilization. A single f1-micro is a poor fit for hosting a production application that requires stability.

upvoted 1 times

JoseCloudEng1994 11 months, 3 weeks ago

Reliable Deployments: GKE provides a superior deployment pipeline. Using Kubernetes Deployments offers features like rolling updates, automated health checks, and rollbacks, which are more robust than the instance template updates in a MIG. The container-based workflow is also faster and more flexible for frequent updates.

Why B and D Are Flawed Option B (MIG with VM Images): While a common approach, building and deploying new VM images for every minor change is a slower and less agile process compared to building and pushing a new container image. It's a less efficient deployment model.

Option D (GKE with Bad Practices): This option is the worst of the bunch. Using an n1-standard-4 machine is massive over-provisioning, completely failing to meet the goal of maximizing utilization. The use of the mutable latest tag and imagePullPolicy: Always is a major anti-pattern in production, leading to unpredictable and irreproducible deployments.

upvoted 1 times

JoseCloudEng1994 Most Recent 11 months, 3 weeks ago

1. Inefficient Resource Utilization (f1-micro): The f1-micro machine type is a shared-core VM, meaning its CPU is a fraction of a physical core and offers burstable performance. The application needs a minimum of 0.1 CPU cores. While an f1-micro can provide this, it is not designed for sustained production workloads and can be unreliable under load. This machine type is best suited for low-traffic websites or small-scale testing, not a production application that needs to maximize utilization. A single f1-micro is a poor fit for hosting a production application that requires stability.

upvoted 1 times

rajrocks171 1 year, 2 months ago

Selected Answer: C

C is correct. Efficient resource usage, GKE schedules multiple pods per node maximizing CPU and Memory utilization. Ideal for small app like in question. GKE provides Reliable deployment with rolling updates, health checks and rollbacks. C also supports safe promotion from staging to prod, version control and automation. Why not A, B, D A - Uses f1-micro instances, which are shared-core, less stable under load, and uses startup scripts which are brittle and hard to manage for versioning. B-

upvoted 4 times

rajrocks171 1 year, 2 months ago

C is correct. Efficient resource usage, GKE schedules multiple pods per node maximizing CPU and Memory utilization. Ideal for small app like in question. GKE provides Reliable deployment with rolling updates, health checks and rollbacks. C also supports safe promotion from staging to prod, version control and automation. Why not A, B, D A - Uses f1-micro instances, which are shared-core, less stable under load, and uses startup scripts which are brittle and hard to manage for versioning. B-B- Requires building and deploying new VM images each time — slow and less flexible for frequent deployments. D - Uses latest tag and imagePullPolicy: Always — this can lead to unpredictable deployments and caching issues. No staging step adds risk.

upvoted 4 times

rajrocks171 1 year, 2 months ago

B- Requires building and deploying new VM images each time — slow and less flexible for frequent deployments. D - Uses latest tag and imagePullPolicy: Always — this can lead to unpredictable deployments and caching issues. No staging step adds risk.

upvoted 3 times

lokiinaction 1 year, 2 months ago

f1-micro fits better for this tiny webapplication

upvoted 1 times

Gayathri1608 1 year, 8 months ago

i think suitable for the given scenario

upvoted 3 times

nareshthumma 1 year, 9 months ago

Answer is C

upvoted 2 times

44fa527 1 year, 11 months ago

should be option C because if you are working in real world, GKE is the best solution for such a case. Furthermore, its reliable, scalable, flexible, at least the best option among the other three.

upvoted 3 times

cai_engineer 1 year, 12 months ago

Ngl it's A. Don't use GKE, it won't schedule the deployment as most of the resources already occupied by kube-systemAlso you can deploy COS Containerd in a VM

upvoted 1 times

cai_engineer 1 year, 12 months ago

Selected Answer: A

Ngl it's A. Don't use GKE, it won't schedule the deployment as most of the resources already occupied by kube-system

upvoted 1 times

cai_engineer 1 year, 12 months ago

Also you can deploy COS Containerd in a VM

upvoted 1 times

awsgcparch 2 years ago

imagePullPolicy: Always ensures that the latest version of the image is always pulled, which guarantees that the most recent code is deployed. Restarting pods ensures that the new version is deployed without requiring manual intervention.

upvoted 1 times

krokskan 2 years, 5 months ago

B because Kubernetes will be overkill and A is not reliable

upvoted 2 times

Gall 2 years, 6 months ago

A is wrong as after the restart the script will be rerun and fetch the code directly from the repo (even if production). The load of the massive number of dependencies will take a lot of timee, and the application version will be fuzzy.

upvoted 3 times

moumou 2 years, 6 months ago

C is correct, B (instance template cannote be updated once created.

upvoted 2 times

kip21 2 years, 7 months ago

C - Correct

upvoted 1 times

AWS_Sam 2 years, 7 months ago

The correct answer is C. Because it is the only option that RELIABLY tests the app in staging before it is applied to production. Remember that one of the requirements in the question is to reliably deploy the app.

upvoted 3 times

Roro_Brother 2 years, 8 months ago

You don't need GKE for 0.1 CPU, only A meet hte needs

upvoted 3 times

MahAli 2 years, 8 months ago

For 0.1 CPU I will never use GKE, considering the associated cost with control plane and not even one option in the question mentioning micro instances for the node pool

upvoted 5 times

Question #68 *Topic 1*

Your company wants to start using Google Cloud resources but wants to retain their on-premises Active Directory domain controller for identity management.
What should you do?

A. Use the Admin Directory API to authenticate against the Active Directory domain controller.

B. Use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML SSO.

C. Use Cloud Identity-Aware Proxy configured to use the on-premises Active Directory domain controller as an identity provider.

D. Use Compute Engine to create an Active Directory (AD) domain controller that is a replica of the on-premises AD domain controller using Google Cloud Directory Sync.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (88%) D (12%)

KouShikyou Highly Voted 6 years, 10 months ago
According to the reference, my understanding is B is correct. And in the document(https://cloud.google.com/iap/docs/concepts-overview), it says: If you need to create Google Accounts for your existing users, you can use Google Cloud Directory Sync to synchronize with your Active Directory or LDAP server.

Is it possible to explain why correct answer is C?
upvoted 45 times

KouShikyou Highly Voted 6 years, 10 months ago
According to the reference, my understanding is B is correct. And in the document(https://cloud.google.com/iap/docs/concepts-overview), it says: If you need to create Google Accounts for your existing users, you can use Google Cloud Directory Sync to synchronize with your Active Directory or LDAP server. Is it possible to explain why correct answer is C?B is okB should be correctB, use GCDS.agreed
upvoted 45 times

tartar 6 years ago
B is ok
upvoted 9 times

kumarp6 5 years, 9 months ago
B should be correct
upvoted 5 times

nitinz 5 years, 5 months ago
B, use GCDS.
upvoted 5 times

AzureDP900 3 years, 10 months ago
agreed
upvoted 1 times

MeasService 6 years, 10 months ago
B is the nearest answer I feel !
upvoted 25 times

eff12c1 Most Recent 2 years, 2 months ago
To integrate Google Cloud with your on-premises Active Directory (AD) domain controller for identity management while retaining your on-premises AD, the best approach is: B. Use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML SSO.
upvoted 1 times

svkds 2 years, 3 months ago
The most suitable option for integrating Google Cloud resources with an on-premises Active Directory domain controller for identity management is option D. This involves creating a replica of the on-premises Active Directory domain controller using Compute Engine and Google Cloud Directory

Sync for synchronization.

upvoted 1 times

LaxmanTiwari 3 years, 3 months ago

B is correct https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-introduction

upvoted 3 times

vamgcp 3 years, 6 months ago

Connect your on-premises Active Directory to Google Cloud: You can use Google Cloud Directory Sync (GCDS) to synchronize your on-premises Active Directory with Google Cloud. This allows you to use your existing Active Directory users and groups in Google Cloud. Set up single sign-on (SSO): You can use Google Cloud Identity-Aware Proxy (IAP) to set up SSO for your Google Cloud resources. IAP integrates with your on-premises Active Directory and allows users to log in to Google Cloud using their existing Active Directory credentials.

upvoted 3 times

omermahgoub 3 years, 8 months ago

B. Use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML SSO.

To retain their on-premises Active Directory domain controller for identity management while using Google Cloud resources, the company can use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML single sign-on (SSO). This will allow users to use their existing Active Directory credentials to access Google Cloud resources, while still maintaining their on-premises Active Directory domain controller as the primary source of identity management.

upvoted 8 times

omermahgoub 3 years, 8 months ago

B. Use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML SSO. To retain their on-premises Active Directory domain controller for identity management while using Google Cloud resources, the company can use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML single sign-on (SSO). This will allow users to use their existing Active Directory credentials to access Google Cloud resources, while still maintaining their on-premises Active Directory domain controller as the primary source of identity management.Option A, using the Admin Directory API to authenticate against the Active Directory domain controller, would not be a suitable solution because it would require implementing custom authentication logic in the application, which would be time-consuming and error-prone. Option C, using Cloud Identity-Aware Proxy configured to use the on-premises Active Directory domain controller as an identity provider, would be a suitable solution, but it would not allow you to synchronize Active Directory usernames with cloud identities. Option D, using Compute Engine to create an Active Directory (AD) domain controller that is a replica of the on-premises AD domain controller using Google Cloud Directory Sync, would not be a suitable solution because it would require setting up and maintaining an additional AD domain controller in Google Cloud, which would be unnecessary if the company wants to retain their on-premises AD domain controller as the primary source of identity management.

upvoted 8 times

omermahgoub 3 years, 8 months ago

Option A, using the Admin Directory API to authenticate against the Active Directory domain controller, would not be a suitable solution because it would require implementing custom authentication logic in the application, which would be time-consuming and error-prone. Option C, using Cloud Identity-Aware Proxy configured to use the on-premises Active Directory domain controller as an identity provider, would be a suitable solution, but it would not allow you to synchronize Active Directory usernames with cloud identities. Option D, using Compute Engine to create an Active Directory (AD) domain controller that is a replica of the on-premises AD domain controller using Google Cloud Directory Sync, would not be a suitable solution because it would require setting up and maintaining an additional AD domain controller in Google Cloud, which would be unnecessary if the company wants to retain their on-premises AD domain controller as the primary source of identity management.

upvoted 3 times

SureshbabuK 3 years, 8 months ago

GCDS and Cloud Identity is provided exactly for this use case

upvoted 1 times

megumin 3 years, 9 months ago

B is ok

upvoted 2 times

Mahmoud_E 3 years, 10 months ago

B is correct https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-introduction

upvoted 2 times

cbarg 4 years, 1 month ago

Answer is B

upvoted 1 times

SAMBIT 4 years, 5 months ago

https://support.google.com/a/answer/106368?hl=en

upvoted 1 times

haroldbenites 4 years, 8 months ago

Go for B. Cloud Directory Sync https://cloud.google.com/blog/products/identity-security/using-your-existing-identity-management-system-with-google-cloud-platform

upvoted 4 times

vincy2202 4 years, 8 months ago

B is the correct answer

upvoted 1 times

pulkit0627 4 years, 9 months ago

B as AD groups are directly mapped to Cloud Directory Sync

upvoted 1 times

MaxNRG 4 years, 9 months ago

B – use Google Cloud Directory Sync to sync Active Directory user names with cloud identities and configure SAML SSO. Check the flowchart here illustrating integration of your existing identity management system into GCP: https://cloud.google.com/blog/products/identity-security/using-your-existing-identity-management-system-with-google-cloud-platform C – does not work, since Cloud IAP serves different purpose. It s a building block toward BeyondCorp, an enterprise security model that enables every employee to work from untrusted networks without the use of a VPN.

upvoted 2 times

MamthaSJ 5 years, 1 month ago

Answer is B

upvoted 4 times

Question #69 *Topic 1*

You are running a cluster on Kubernetes Engine (GKE) to serve a web application. Users are reporting that a specific part of the application is not responding anymore. You notice that all pods of your deployment keep restarting after 2 seconds. The application writes logs to standard output. You want to inspect the logs to find the cause of the issue. Which approach can you take?

A. Review the Stackdriver logs for each Compute Engine instance that is serving as a node in the cluster.

B. Review the Stackdriver logs for the specific GKE container that is serving the unresponsive part of the application.

C. Connect to the cluster using gcloud credentials and connect to a container in one of the pods to read the logs.

D. Review the Serial Port logs for each Compute Engine instance that is serving as a node in the cluster.

Suggested Answer: *B*

Official Answer: B

Community vote distribution

B (100%)

jcmoranp Highly Voted 5 years, 10 months ago

think answer is B. C cannot be, you don't need to connect to the container to view logs, you connect to stackdriver for this

upvoted 34 times

jcmoranp Highly Voted 5 years, 10 months ago

think answer is B. C cannot be, you don't need to connect to the container to view logs, you connect to stackdriver for thisIs Stackdriver enabled by default? Stackdriver Logging is independent and first needs to enable with GKE I guess?B is okPlease forget this comment ^ Answer B should be correct.Yes for GKEYes it is BB, google wants you to use stackdriver.Stackdriver Logging seems to be enabled by default for GKE. Looking here: https://cloud.google.com/monitoring/kubernetes-engine/legacy-stackdriver/logging For container and system logs, GKE deploys a per-node logging agent that reads container logs, adds helpful metadata, and then stores them. The logging agent checks for container logs in the following sources: Standard output and standard error logs from containerized processes I would also go with Bagreed with B

upvoted 34 times

crypt0 5 years, 10 months ago

Is Stackdriver enabled by default? Stackdriver Logging is independent and first needs to enable with GKE I guess?

upvoted 1 times

crypt0 5 years, 10 months ago

Is Stackdriver enabled by default? Stackdriver Logging is independent and first needs to enable with GKE I guess?B is okPlease forget this comment ^ Answer B should be correct.Yes for GKE

upvoted 1 times

tartar 5 years ago

B is ok

upvoted 7 times

crypt0 5 years, 9 months ago

Please forget this comment ^ Answer B should be correct.

upvoted 8 times

Jack_in_Large 5 years, 3 months ago

Yes for GKE

upvoted 1 times

kumarp6 4 years, 9 months ago

Yes it is B

upvoted 3 times

nitinz 4 years, 5 months ago

B, google wants you to use stackdriver.

upvoted 6 times

crypt0 5 years, 9 months ago

Stackdriver Logging seems to be enabled by default for GKE.

Looking here: https://cloud.google.com/monitoring/kubernetes-engine/legacy-stackdriver/logging For container and system logs, GKE deploys a per-node logging agent that reads container logs, adds helpful metadata, and then stores them. The logging agent checks for container logs in the following sources:

Standard output and standard error logs from containerized processes

I would also go with B

upvoted 13 times

crypt0 5 years, 9 months ago

Stackdriver Logging seems to be enabled by default for GKE. Looking here: https://cloud.google.com/monitoring/kubernetes-engine/legacy-stackdriver/logging For container and system logs, GKE deploys a per-node logging agent that reads container logs, adds helpful metadata, and then stores them. The logging agent checks for container logs in the following sources: Standard output and standard error logs from containerized processes I would also go with Bagreed with B

upvoted 13 times

AzureDP900 2 years, 10 months ago

agreed with B

upvoted 1 times

JoeShmoe 5 years, 9 months ago

B is correct. Serial console doesnt give you StdOut

upvoted 9 times

AdityaGupta Most Recent 1 year, 10 months ago

B. Review the Stackdriver logs for the specific GKE container that is serving the unresponsive part of the application. GKE be default integrats with Google Operation Suit (Stackdriver) and you can filter the logs for more specific part of application i.e container to view logs. Also it is most efficient way of investigation.

upvoted 2 times

jalberto 2 years ago

B is the simples option and more effective

upvoted 1 times

MestreCholas 2 years, 5 months ago

B. Review the Stackdriver logs for the specific GKE container that is serving the unresponsive part of the application. Since the application writes logs to standard output, the logs should be available in the Stackdriver logs for the container running the unresponsive part of the application. Kubernetes Engine automatically exports these logs to Stackdriver, so you can use the Stackdriver Logging console to view the logs. Option A is not the best choice because reviewing the logs for each Compute Engine instance would be time-consuming and may not provide the necessary information. Option C may work, but it involves extra steps and may not be necessary if the logs are available in Stackdriver. Option D is not relevant in this case because Serial Port logs are not likely to provide useful information for troubleshooting an unresponsive web application.

upvoted 1 times

omermahgoub 2 years, 8 months ago

C. Connect to the cluster using gcloud credentials and connect to a container in one of the pods to read the logs.

To inspect the logs of a Kubernetes Engine (GKE) cluster to find the cause of an issue, you can connect to the cluster using gcloud credentials and connect to a container in one of the pods to read the logs. This will allow you to access the logs of the application as it is running in the cluster, which should help you identify the cause of the issue.

upvoted 2 times

omermahgoub 2 years, 8 months ago

C. Connect to the cluster using gcloud credentials and connect to a container in one of the pods to read the logs. To inspect the logs of a Kubernetes Engine (GKE) cluster to find the cause of an issue, you can connect to the cluster using gcloud credentials and connect to a container in one of the pods to read the logs. This will allow you to access the logs of the application as it is running in the cluster, which should help you identify the cause of the issue.Option A, reviewing the Stackdriver logs for each Compute Engine instance that is serving as a node in the cluster, would not be suitable because the application writes logs to standard output, not to Stackdriver. Option B, reviewing the Stackdriver logs for the specific GKE container that is serving the unresponsive part of the application, would not be suitable because the application writes logs to standard output, not to Stackdriver. Option D, reviewing the Serial Port logs for each Compute Engine instance that is serving as a node in the cluster, would not be suitable because the application writes logs to standard output, not to the Serial Port.

upvoted 2 times

omermahgoub 2 years, 8 months ago

Option A, reviewing the Stackdriver logs for each Compute Engine instance that is serving as a node in the cluster, would not be suitable because the application writes logs to standard output, not to Stackdriver. Option B, reviewing the Stackdriver logs for the specific GKE container that is serving the unresponsive part of the application, would not be suitable because the application writes logs to standard output, not to Stackdriver. Option D, reviewing the Serial Port logs for each Compute Engine instance that is serving as a node in the cluster, would not be suitable because the application writes logs to standard output, not to the Serial Port.

upvoted 3 times

jaxclain 2 years, 8 months ago

This should be easy, the answer is B. Just eliminate the wrong answers (A) is not correct because the question is about GKE and not CE. C and D are totally lost

upvoted 1 times

habros 2 years, 8 months ago

A, C, D all sounds unfeasible (credentials, and compute engine)

upvoted 1 times

megumin 2 years, 9 months ago

B is ok

upvoted 1 times

Mahmoud_E 2 years, 10 months ago

is correct answer

upvoted 1 times

Superr 3 years, 2 months ago

correct

upvoted 1 times

OrangeTiger 3 years, 7 months ago

B is correct ans.I agree. https://cloud.google.com/blog/ja/products/management-tools/finding-your-gke-logs

upvoted 2 times

haroldbenites 3 years, 8 months ago

Go for B

upvoted 1 times

vincy2202 3 years, 8 months ago

B is the correct answer

upvoted 1 times

MaxNRG 3 years, 9 months ago

B – Review Stackdriver logs for specific GKE container that is serving the unresponsive part of the app. This is a most directly matching answer for this Q, since it reviews GKE container logs, by that advertising this Stackdriver feature. “For container and system logs, GKE deploys a per-node logging agent that reads container logs, adds helpful metadata, and then stores them. The logging agent checks for container logs in the following sources: • Standard output and standard error logs from containerized processes • kubelet and container runtime logs • Logs for system components, such as VM startup scripts” Originally we thought, that D is a right answer, since were confused with 2 seconds restart. But, that’s restart for Pod, not for Node (GCE) D – Review Serial Port logs for each Compute Enginee instance, that is serving as the in the cluster. Serial Port output is standard feature of Compute Engine (which retains 1 MB most recent logs for analysis). But, it is irrelevant for Pod’s restart, caused by malfunction of some container.

upvoted 3 times

MamthaSJ 4 years, 1 month ago

Answer is B

upvoted 3 times

lovingsmart2000 4 years, 1 month ago

B is right answer. There is a catch here - Legacy logging of GKE with Stackdriver has deprecated. If this is used, you need to migrate to Cloud Operations for GKE, a new enhanced offering by Google with same functionality. Future questions will have the answer choices with new tool "Cloud Operations for GKE" instead of Stackdriver. https://cloud.google.com/monitoring/kubernetes-engine/legacy-stackdriver/logging

upvoted 2 times

Question #70 *Topic 1*

You are using a single Cloud SQL instance to serve your application from a specific zone. You want to introduce high availability. What should you do?

A. Create a read replica instance in a different region

B. Create a failover replica instance in a different region

C. Create a read replica instance in the same region, but in a different zone

D. Create a failover replica instance in the same region, but in a different zone

Suggested Answer: *D*

Official Answer: D

Community vote distribution

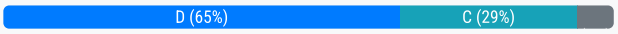

AWS56 Highly Voted 6 years, 7 months ago

Agree D

upvoted 36 times

AWS56 Highly Voted 6 years, 7 months ago

Agree DD is okYes D is rightthis Question is very Old and should be deleted from the exam , there is no Failover replica now , to do an HA we just confer it for the SQL instance that we have .https://cloud.google.com/sql/docs/mysql/replication#:~:text=Read%20replicas%20neither%20provide%20high%20availability%20nor%20offer%20it.&text=A%20primary%20instance%20cannot%20failover,any%20way%20during%20an%20outage.&text=Maintenance%20windows%20cannot%20be%20set,windows%20with%20the%20primary%20instance. - Read replicas neither provide high availability nor offer it. Agree DThat link is helpful! I navigated to the "quick reference for Cloud SQL read replicas" and read the "failover" and "high availability" topics. They state: 1. Failover - "A primary instance cannot failover to a read replica, and read replicas are unable to failover in any way during an outage." 2. High Availability - "Read replicas neither provide high availability nor offer it."Sorry for this, but in the same link you can read: High availability Read replicas allow you to enable high availability on the replicas. (What I understand is that Read Replicas give you high availability on reads, of course, not in writes).

upvoted 36 times

tartar 6 years ago

D is ok

upvoted 7 times

kumarp6 5 years, 9 months ago

Yes D is right

upvoted 4 times

kimharsh 4 years, 2 months ago

this Question is very Old and should be deleted from the exam , there is no Failover replica now , to do an HA we just confer it for the SQL instance that we have .

upvoted 25 times

nwk 4 years, 2 months ago

https://cloud.google.com/sql/docs/mysql/replication#:~:text=Read%20replicas%20neither%20provide%20high%20availability%20nor%20offer%20it.&text=A%20primary%20instance%20cannot%20failover,any%20way%20during%20an%20outage.&text=Maintenance%20windows%20cannot%20be%20set,windows%20with%20the%20primary%20instance. - Read replicas neither provide high availability nor offer it. Agree DThat link is helpful! I navigated to the "quick reference for Cloud SQL read replicas" and read the "failover" and "high availability" topics. They state: 1. Failover - "A primary instance cannot failover to a read replica, and read replicas are unable to failover in any way during an outage." 2. High Availability - "Read replicas neither provide high availability nor offer it."Sorry for this, but in the same link you can read: High availability Read replicas allow you to enable high availability on the replicas. (What I understand is that Read Replicas give you high availability on reads, of course, not in writes).

upvoted 7 times

nwk 4 years, 2 months ago

https://cloud.google.com/sql/docs/mysql/replication#:~:text=Read%20replicas%20neither%20provide%20high%20availability%20nor%20offer%20it.&text=A%20primary%20instance%20cannot%20failover,any%20way%20during%20an%20outage.&text=Maintenance%20windows%20cannot%20be%20set,windows%20with%20the%20primary%20instance.

- Read replicas neither provide high availability nor offer it.

Agree D

upvoted 7 times

jay9114 3 years, 11 months ago

That link is helpful! I navigated to the "quick reference for Cloud SQL read replicas" and read the "failover" and "high availability" topics. They state: 1. Failover - "A primary instance cannot failover to a read replica, and read replicas are unable to failover in any way during an outage." 2. High Availability - "Read replicas neither provide high availability nor offer it."

upvoted 2 times

jay9114 3 years, 11 months ago

That link is helpful! I navigated to the "quick reference for Cloud SQL read replicas" and read the "failover" and "high availability" topics. They state: 1. Failover - "A primary instance cannot failover to a read replica, and read replicas are unable to failover in any way during an outage." 2. High Availability - "Read replicas neither provide high availability nor offer it."Sorry for this, but in the same link you can read: High availability Read replicas allow you to enable high availability on the replicas. (What I understand is that Read Replicas give you high availability on reads, of course, not in writes).

upvoted 2 times

spuyol 2 years, 6 months ago

Sorry for this, but in the same link you can read: High availability Read replicas allow you to enable high availability on the replicas. (What I understand is that Read Replicas give you high availability on reads, of course, not in writes).

upvoted 1 times

GunjGupta 6 years, 3 months ago

Cloud SQL is regional. For high availability, we need to think fo a failover strategy. So Option D meets the requirement. create failover replica in the same region but in different Zone

upvoted 18 times

ahmed75.zahran00 Most Recent 7 months ago

A read replica is for scaling read traffic, not automatic failover.

upvoted 1 times

JoseCloudEng1994 11 months, 3 weeks ago

The correct answer is D. Create a failover replica instance in the same region, but in a different zone. To provide high availability for a Cloud SQL instance, the standard practice is to use a failover replica. A failover replica is a standby instance that is kept in sync with the primary instance and is located in a different zone within the same region. This configuration is known as a regional instance and provides automatic failover.A. Create a read replica instance in a different region: This is incorrect. A read replica is used for offloading read traffic to improve performance, not for providing high availability with automatic failover. While you can manually promote a read replica to a primary instance, this is a disaster recovery strategy, not a seamless high-availability solution. Furthermore, placing it in a different region addresses regional outages, which is a different use case from protecting against a single zone outage. B. Create a failover replica instance in a different region: This is incorrect. Cloud SQL's high availability feature for automatic failover works within a single region, not across regions. Cross-region setups are typically used for disaster recovery (DR), not for high availability within a region.C. Create a read replica instance in the same region, but in a different zone: This is incorrect. A read replica is not a failover replica. As mentioned above, it is for read-heavy workloads and does not support automatic failover. D. Create a failover replica instance in the same region, but in a different zone: This is the correct approach. Cloud SQL's high availability configuration automatically creates a standby instance (failover replica) in a different zone from the primary instance. If the primary instance or its zone fails, Cloud SQL automatically switches over to the standby instance. The standby instance takes over the primary's static IP address, providing a seamless failover with minimal downtime.

upvoted 2 times

JoseCloudEng1994 11 months, 3 weeks ago

Selected Answer: D

The correct answer is D. Create a failover replica instance in the same region, but in a different zone.

To provide high availability for a Cloud SQL instance, the standard practice is to use a failover replica. A failover replica is a standby instance that is kept in sync with the primary instance and is located in a different zone within the same region. This configuration is known as a regional instance and provides automatic failover.

upvoted 2 times

JoseCloudEng1994 11 months, 3 weeks ago

A. Create a read replica instance in a different region: This is incorrect. A read replica is used for offloading read traffic to improve performance, not for providing high availability with automatic failover. While you can manually promote a read replica to a primary instance, this is a disaster recovery strategy, not a seamless high-availability solution. Furthermore, placing it in a different region addresses regional outages, which is a different use

case from protecting against a single zone outage. B. Create a failover replica instance in a different region: This is incorrect. Cloud SQL's high availability feature for automatic failover works within a single region, not across regions. Cross-region setups are typically used for disaster recovery (DR), not for high availability within a region.

upvoted 1 times

JoseCloudEng1994 11 months, 3 weeks ago

A. Create a read replica instance in a different region: This is incorrect. A read replica is used for offloading read traffic to improve performance, not for providing high availability with automatic failover. While you can manually promote a read replica to a primary instance, this is a disaster recovery strategy, not a seamless high-availability solution. Furthermore, placing it in a different region addresses regional outages, which is a different use case from protecting against a single zone outage. B. Create a failover replica instance in a different region: This is incorrect. Cloud SQL's high availability feature for automatic failover works within a single region, not across regions. Cross-region setups are typically used for disaster recovery (DR), not for high availability within a region.C. Create a read replica instance in the same region, but in a different zone: This is incorrect. A read replica is not a failover replica. As mentioned above, it is for read-heavy workloads and does not support automatic failover. D. Create a failover replica instance in the same region, but in a different zone: This is the correct approach. Cloud SQL's high availability configuration automatically creates a standby instance (failover replica) in a different zone from the primary instance. If the primary instance or its zone fails, Cloud SQL automatically switches over to the standby instance. The standby instance takes over the primary's static IP address, providing a seamless failover with minimal downtime.

upvoted 1 times

JoseCloudEng1994 11 months, 3 weeks ago

C. Create a read replica instance in the same region, but in a different zone: This is incorrect. A read replica is not a failover replica. As mentioned above, it is for read-heavy workloads and does not support automatic failover. D. Create a failover replica instance in the same region, but in a different zone: This is the correct approach. Cloud SQL's high availability configuration automatically creates a standby instance (failover replica) in a different zone from the primary instance. If the primary instance or its zone fails, Cloud SQL automatically switches over to the standby instance. The standby instance takes over the primary's static IP address, providing a seamless failover with minimal downtime.

upvoted 1 times

matt57 1 year, 6 months ago

D - It's serving completely your application, so it should not be read replica it should be failover replica on a diff zone

upvoted 2 times

plumbig11 1 year, 7 months ago

Read replica is not appropriated to failover. Create a failover replica instance in the same region, but in a different zone

upvoted 2 times

Lenifia 1 year, 9 months ago

C is the correct answer, failover is more like DR

upvoted 1 times

nareshthumma 1 year, 9 months ago

Answer is D: Create a failover replica instance in the same region, but in a different zone. Here's why this option is the best: 1. High Availability: A failover replica provides automatic failover capabilities, meaning that if the primary instance becomes unavailable (due to a zone failure, for example), Cloud SQL can automatically promote the failover replica to be the new primary instance, minimizing downtime. 2. Same Region, Different Zone: By creating the failover replica in the same region but a different zone, you ensure that the instances are geographically close to each other, which helps maintain low latency and faster failover times while still protecting against zone-specific outages. 3. Cost Efficiency: Using a failover replica in the same region is typically more cost-effective than setting up a replica in a different region, as cross-region replication can introduce additional latency and costs.

upvoted 2 times

hehe_24 1 year, 10 months ago

C is incorrect. D is the right answer. Read-replica is just an instance for read operation, it cannot provide HA.

upvoted 1 times

0verK0alafied 2 years, 3 months ago

The HA configuration provides data redundancy. A Cloud SQL instance configured for HA is also called a regional instance and has a primary and secondary zone within the configured region. Within a regional instance, the configuration is made up of a primary instance and a standby instance. Through synchronous replication to each zone's persistent disk, all writes made to the primary instance are replicated to disks in both zones before a transaction is reported as committed. In the event of an instance or zone failure, the standby instance becomes the new primary instance. Users are then rerouted to the new primary instance. This process is called a failover. https://cloud.google.com/sql/docs/mysql/high-availability#HA-configuration

upvoted 4 times

Gall 2 years, 6 months ago

"Note: Read replicas do not provide failover capability. To provide failover capability for an instance, see Configuring an instance for high availability." https://cloud.google.com/sql/docs/mysql/replication/

upvoted 3 times

parthkulkarni998 2 years, 8 months ago

Correct answer would be D as a failover replica acts as a redundant copy incase of zone failure. However, option C causes confusion because a read replica can provide availability for reads, in case of zone failure for primary, but they cant provide support for writes. They would only work for reads.

upvoted 2 times

Roro_Brother 2 years, 8 months ago

D, its regional product and failover is required for HA

upvoted 2 times

thewalker 2 years, 9 months ago

C 1. Failover replica is a legacy way and is not available in GCP now - B and D are not the options: https://cloud.google.com/sql/docs/mysql/high-availability#legacy_mysql_high_availability_option 2. Cloud SQL is regional resource. However, cross-region read replicas are allowed now in Cloud SQL (https://cloud.google.com/blog/products/databases/introducing-cross-region-replica-for-cloud-sql) - A and C are options. Chosen C, as there is no requirement or mention of cross-regional / global db.

upvoted 1 times

thewalker 2 years, 9 months ago

C 1. Failover replica is a legacy way and is not available in GCP now - B and D are not the options: https://cloud.google.com/sql/docs/mysql/high-availability#legacy_mysql_high_availability_option 2. Cloud SQL is regional resource. However, cross-region read replicas are allowed now in Cloud SQL (https://cloud.google.com/blog/products/databases/introducing-cross-region-replica-for-cloud-sql) - A and C are options. Chosen C, as there is no requirement or mention of cross-regional / global db.Read replica is not a valid choice for HA configurations. It does not provide automatic failover that is required for HA. It may be called something different or this answer has changed, but D is still the best option.

upvoted 1 times

hogtrough 2 years, 7 months ago

Read replica is not a valid choice for HA configurations. It does not provide automatic failover that is required for HA. It may be called something different or this answer has changed, but D is still the best option.

upvoted 1 times

piiizu 2 years, 10 months ago

The key is High Availability, not Resilience or Disaster Recovery. Therefore my answer is C

upvoted 1 times

someone2011 2 years, 11 months ago

D. In HA config, the second replica is caled stand by. The process of replacing the primary damaged node is called failover. https://cloud.google.com/sql/docs/postgres/high-availability

upvoted 2 times

didek1986 2 years, 11 months ago

C for sure

upvoted 1 times

Question #71 *Topic 1*

Your company is running a stateless application on a Compute Engine instance. The application is used heavily during regular business hours and lightly outside of business hours. Users are reporting that the application is slow during peak hours. You need to optimize the application's performance. What should you do?

A. Create a snapshot of the existing disk. Create an instance template from the snapshot. Create an autoscaled managed instance group from the instance template.

B. Create a snapshot of the existing disk. Create a custom image from the snapshot. Create an autoscaled managed instance group from the custom image.

C. Create a custom image from the existing disk. Create an instance template from the custom image. Create an autoscaled managed instance group from the instance template.

D. Create an instance template from the existing disk. Create a custom image from the instance template. Create an autoscaled managed instance group from the custom image.

Suggested Answer: *C*
Official Answer: C

Community vote distribution
C (95%) (5%)

sdsdfasdf4 Highly Voted 5 years, 8 months ago
The easiest way would be to create template from --source-instance, and then create MIG, but it is not listed here, also you cannot create a MIG from image directly, you need a template, so answer is C (image -> template -> mig).
upvoted 32 times

sdsdfasdf4 Highly Voted 5 years, 8 months ago
The easiest way would be to create template from --source-instance, and then create MIG, but it is not listed here, also you cannot create a MIG from image directly, you need a template, so answer is C (image -> template -> mig).C is correct. To sdsdfasd4's point - Not recommended to create template from --source-instance as If the existing instance contains a static external IP address, that address is copied into the instance template and might limit the use of the template. Templates are best created from images or other templates. Creating the template from a running instance may require work to clean it up before it can be used for a MIG
upvoted 32 times

6721sora 3 years, 11 months ago
C is correct. To sdsdfasd4's point - Not recommended to create template from --source-instance as If the existing instance contains a static external IP address, that address is copied into the instance template and might limit the use of the template. Templates are best created from images or other templates. Creating the template from a running instance may require work to clean it up before it can be used for a MIG
upvoted 9 times

AWS56 6 years, 7 months ago
C is the right answer
upvoted 12 times

AWS56 6 years, 7 months ago
C is the right answerC is definitely the right answer
upvoted 12 times

heretolearnazure 2 years, 12 months ago
C is definitely the right answer
upvoted 1 times

nareshthumma Most Recent 1 year, 9 months ago
Answer A: Create a snapshot of the existing disk. Create an instance template from the snapshot. Create an autoscaled managed instance group from the instance template. Here's why this option is the best: 1. Autoscaling: By creating an autoscaled managed instance group, you can automatically adjust the number of instances based on the load. This means that during peak business hours, additional instances will be created to handle the increased traffic, improving application performance. During off-peak hours, the number of instances can scale down to reduce costs. 2. Snapshot Creation: Taking a snapshot of the existing disk ensures that you have a backup of the current state of your application. This snapshot can be used to create the new instance template, ensuring that your autoscaled instances have the same configuration as the original instance.

upvoted 1 times

46f094c 2 years, 1 month ago

I prefer B, is better because I don't need to stop the instance to create the disk image.

upvoted 1 times

pakilodi 2 years, 8 months ago

C. The key here is stateless, so we don't need snapshot the actual instance, we can start from zero.

upvoted 4 times

thewalker 2 years, 8 months ago

C We can create an instance from an image or a custom image or a snapshot but an instance template can be created using either image or custom image only. Also refer: https://cloud.google.com/compute/docs/instance-templates/create-instance-templates

upvoted 4 times

Deb2293 3 years, 5 months ago

Option C is the correct choice because creating a custom image from the existing disk ensures that the application environment is consistent and does not change between instances, which can reduce variability in performance. Creating an instance template from the custom image allows you to easily create new instances that are based on the same image, which can save time and effort. Finally, creating an autoscaled managed instance group allows you to automatically scale the number of instances based on demand, which can ensure that there are enough instances to handle peak traffic while minimizing costs during periods of low traffic

upvoted 7 times

megumin 3 years, 9 months ago

C is ok

upvoted 2 times

AzureDP900 3 years, 10 months ago

C is right. Create a custom image from the existing disk. Create an instance template from the custom image. Create an autoscaled managed instance group from the instance template.

upvoted 1 times

abirroy 3 years, 11 months ago

C is correct

upvoted 1 times

mv2000 4 years, 1 month ago

06/30/2022 Exam

upvoted 9 times

SHalgatti 4 years, 6 months ago

I think Snapshot option are not correct in this scenario as to take snapshot you need to stop the VM so C looks best option

upvoted 2 times

SHalgatti 4 years, 6 months ago

I think Snapshot option are not correct in this scenario as to take snapshot you need to stop the VM so C looks best optionit's possible to create a snapshot of running VM by reducing I/O disks

upvoted 2 times

bargou 2 years, 6 months ago

it's possible to create a snapshot of running VM by reducing I/O disks

upvoted 1 times

PhuocT 4 years, 7 months ago

C is the answer

upvoted 1 times

Rajasa 4 years, 7 months ago

C is the answer

upvoted 1 times

haroldbenites 4 years, 8 months ago

Go for C. Instance template can not be created from snapshot. Only from an image.

upvoted 3 times

vincy2202 4 years, 8 months ago

C is the right answer. https://cloud.google.com/compute/docs/instance-templates/create-instance-templates#using_custom_or_public_images_in_your_instance_templates

upvoted 3 times

joe2211 4 years, 8 months ago

vote C

upvoted 1 times

Question #72 *Topic 1*

Your web application has several VM instances running within a VPC. You want to restrict communications between instances to only the paths and ports you authorize, but you don't want to rely on static IP addresses or subnets because the app can autoscale. How should you restrict communications?

A. Use separate VPCs to restrict traffic

B. Use firewall rules based on network tags attached to the compute instances

C. Use Cloud DNS and only allow connections from authorized hostnames

D. Use service accounts and configure the web application to authorize particular service accounts to have access

Suggested Answer: *B*

Official Answer: B

Community vote distribution

B (100%)

AWS56 Highly Voted 5 years, 1 month ago

Agree BYes B it isB is correct

upvoted 24 times

AWS56 Highly Voted 5 years, 1 month ago

Agree B

upvoted 24 times

kumarp6 4 years, 3 months ago

Yes B it is

upvoted 2 times

nitinz 3 years, 11 months ago

B is correct

upvoted 2 times

omermahgoub 2 years, 1 month ago

B. Use firewall rules based on network tags attached to the compute instances

To restrict communications between VM instances within a VPC without relying on static IP addresses or subnets, you can use firewall rules based on network tags attached to the compute instances. This will allow you to specify which instances are allowed to communicate with each other and on which paths and ports. You can then attach the relevant network tags to the compute instances when they are created, allowing you to control communication between the instances without relying on static IP addresses or subnets.

upvoted 13 times

omermahgoub 2 years, 1 month ago

B. Use firewall rules based on network tags attached to the compute instances To restrict communications between VM instances within a VPC without relying on static IP addresses or subnets, you can use firewall rules based on network tags attached to the compute instances. This will allow you to specify which instances are allowed to communicate with each other and on which paths and ports. You can then attach the relevant network tags to the compute instances when they are created, allowing you to control communication between the instances without relying on static IP addresses or subnets.Option A, using separate VPCs to restrict traffic, would not be a suitable solution because it would not allow the instances to communicate with each other, which is likely necessary for the functioning of the web application. Option C, using Cloud DNS and only allowing connections from authorized hostnames, would not be a suitable solution because it would not allow you to control communication between the instances based on their IP addresses or other characteristics. Option D, using service accounts and configuring the web application to authorize particular service accounts to have access, would not be a suitable solution because it would not allow you to control communication between the instances based on their IP addresses or other characteristics.

upvoted 13 times

omermahgoub 2 years, 1 month ago

Option A, using separate VPCs to restrict traffic, would not be a suitable solution because it would not allow the instances to communicate with each other, which is likely necessary for the functioning of the web application. Option C, using Cloud DNS and only allowing connections from authorized hostnames, would not be a suitable solution because it would not allow you to control communication between the instances based on their IP

addresses or other characteristics. Option D, using service accounts and configuring the web application to authorize particular service accounts to have access, would not be a suitable solution because it would not allow you to control communication between the instances based on their IP addresses or other characteristics.

upvoted 5 times

SureshbabuK 2 years, 2 months ago

Access to specific ports and protocol can be controlled only by firewall rule - Hence B is correct. D is not correct as service account is to authenticate and Authorized a specific machine to resource or service not ports and protocols

upvoted 4 times

megumin 2 years, 3 months ago

B is ok

upvoted 1 times

AzureDP900 2 years, 4 months ago

B is the best option.

upvoted 1 times

abirroy 2 years, 5 months ago

Use firewall rules based on network tags attached to the compute instances

upvoted 2 times

alexandercamachop 2 years, 5 months ago

The secret is "paths and ports". Which tell us Firewall as our only option.And how does firewall restrict "paths" pretty please?

upvoted 5 times

alexandercamachop 2 years, 5 months ago

The secret is "paths and ports". Which tell us Firewall as our only option.

upvoted 5 times

medi01 Most Recent 1 year, 10 months ago

And how does firewall restrict "paths" pretty please?

upvoted 5 times

cbarg 2 years, 7 months ago

B Firewall rules to restrict traffic

upvoted 1 times

haroldbenites 3 years, 2 months ago

Go for B.

upvoted 2 times

vincy2202 3 years, 2 months ago

B is the right answer

upvoted 2 times

MaxNRG 3 years, 3 months ago

B – use firewall rules based on network tags attached to the compute instances This answer avoids using IP, which are replaced by tags.

upvoted 3 times

MamthaSJ 3 years, 7 months ago

Answer is B

upvoted 4 times

areza 3 years, 8 months ago

B is ok

upvoted 1 times

victory108 3 years, 9 months ago

B. Use firewall rules based on network tags attached to the compute instances

upvoted 2 times

Ausias18 3 years, 10 months ago

Answer is B

upvoted 1 times

lynx256 3 years, 10 months ago

B is ok

upvoted 1 times

Vika 3 years, 11 months ago

Agree B

upvoted 1 times

Question #73 *Topic 1*

You are using Cloud SQL as the database backend for a large CRM deployment. You want to scale as usage increases and ensure that you don't run out of storage, maintain 75% CPU usage cores, and keep replication lag below 60 seconds. What are the correct steps to meet your requirements?

A. 1. Enable automatic storage increase for the instance. 2. Create a Stackdriver alert when CPU usage exceeds 75%, and change the instance type to reduce CPU usage. 3. Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.

B. 1. Enable automatic storage increase for the instance. 2. Change the instance type to a 32-core machine type to keep CPU usage below 75%. 3. Create a Stackdriver alert for replication lag, and deploy memcache to reduce load on the master.

C. 1. Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space. 2. Deploy memcached to reduce CPU load. 3. Change the instance type to a 32-core machine type to reduce replication lag.

D. 1. Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space. 2. Deploy memcached to reduce CPU load. 3. Create a Stackdriver alert for replication lag, and change the instance type to a 32-core machine type to reduce replication lag.

Suggested Answer: *A*
Official Answer: A

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

AWS56 Highly Voted 6 years, 7 months ago
Agree with A
upvoted 27 times

AWS56 Highly Voted 6 years, 7 months ago
Agree with ASharding database will reduce latency1. Enable automatic storage increase for the instance. 2. Create a Stackdriver alert when CPU usage exceeds 75%, and change the instance type to reduce CPU usage. 3. Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
upvoted 27 times

heretolearnazure 2 years, 12 months ago
Sharding database will reduce latency
upvoted 4 times

AzureDP900 3 years, 10 months ago
1. Enable automatic storage increase for the instance. 2. Create a Stackdriver alert when CPU usage exceeds 75%, and change the instance type to reduce CPU usage. 3. Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
upvoted 2 times

9xnine 4 years, 2 months ago
Has anyone who has taken the exam recently seen any lingering questions with the Stackdriver nomenclature or is it all cloud logging, cloud monitoring, etc.?
upvoted 17 times

Ral17 Most Recent 4 months, 1 week ago
Google's own documentation recommends sharding to reduce replication lag. Also no machine size is mentioned so why do we pick up 32 cores?
upvoted 2 times

nareshthumma 1 year, 9 months ago
Answer is: A
upvoted 1 times

e5019c6 2 years, 7 months ago
Selected Answer: B
While I understand the doubts of selecting a 32 core machine from the start, answer A might be wrong... According to this article: https://cloud.google.com/sql/docs/mysql/instance-settings#impact "For MySQL instances, changing either the machine type or the zone of the instance results in the instance going offline for several minutes." And I understand that instance type == machine type. If we switch to a PostgreSQL or SQL Server instance, similar warnings appear: PostgreSQL: "Changing the number of CPUs or the memory size results in the instance going offline

for less than 60 seconds. The total time for the changes to take effect can take several minutes." SQL Server: "Changing the number of CPUs or the memory size results in the instance going offline for less than 60 seconds."

upvoted 2 times

e5019c6 2 years, 7 months ago

While I understand the doubts of selecting a 32 core machine from the start, answer A might be wrong... According to this article: https://cloud.google.com/sql/docs/mysql/instance-settings#impact "For MySQL instances, changing either the machine type or the zone of the instance results in the instance going offline for several minutes." And I understand that instance type == machine type. If we switch to a PostgreSQL or SQL Server instance, similar warnings appear: PostgreSQL: "Changing the number of CPUs or the memory size results in the instance going offline for less than 60 seconds. The total time for the changes to take effect can take several minutes." SQL Server: "Changing the number of CPUs or the memory size results in the instance going offline for less than 60 seconds."But I also understand that the question didn't mention anything about needing the DB to be online always, so maybe these offline times are acceptable...

upvoted 2 times

e5019c6 2 years, 7 months ago

But I also understand that the question didn't mention anything about needing the DB to be online always, so maybe these offline times are acceptable...

upvoted 1 times

AdityaGupta 2 years, 10 months ago

You are using Cloud SQL as the database backend for a large CRM deployment. You want to scale as usage increases and ensure that you don't run out of storage, maintain 75% CPU usage cores, and keep replication lag below 60 seconds. What are the correct steps to meet your requirements? C & D is out of question as it is talking of 75% of storage, where in question it says 75% of CPU. Option A says monitoring before before taking action and sharding will also help in reducing latency. Option B specifies specific machine type, which is not correct and also memcache which is used to recude the round trip to fetch data, it will help in reducing latency. I would prefer to go with Option A, as it is correct sequence to solve the problem.

upvoted 5 times

red_panda 3 years, 2 months ago

For me is A.

upvoted 1 times

jfricker 3 years, 6 months ago

The correct answer is D. 1. Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space. 2. Deploy memcached to reduce CPU load. 3. Create a Stackdriver alert for replication lag, and change the instance type to a 32-core machine type to reduce replication lag. This approach ensures that you are able to address the three requirements specified in the question: - Monitoring storage usage and increasing storage when it exceeds 75% to avoid running out of storage. - Reducing CPU load by deploying memcached, which can be used to cache frequently-used data, offloading some of the load from the database. - Monitoring replication lag and increasing the number of cores to reduce lag.

upvoted 4 times

Charsoft 3 years, 7 months ago

It may be A for the simple fact that all the other answers throw in a tiny detail about 32 cores. This seems like a red herring (unnecessary details that are meant to distract), so for that reason, A is the answer.

upvoted 5 times

megumin 3 years, 9 months ago

A is ok

upvoted 1 times

AzureDP900 3 years, 10 months ago

A is right

upvoted 1 times

6721sora 3 years, 12 months ago

A is incorrect. because of the wording "Shard the database". How can you shard the database in Cloud SQL without causing major disruptions? Sharding is not a core feature of RDBMS.

B should be correct. inspite of the mention of a fixed 32 core

upvoted 3 times

6721sora 3 years, 12 months ago

A is incorrect. because of the wording "Shard the database". How can you shard the database in Cloud SQL without causing major disruptions? Sharding is not a core feature of RDBMS. B should be correct. inspite of the mention of a fixed 32 coreli Kk Kk You can shard cloudsql. Review this article - https://cloud.google.com/community/tutorials/horizontally-scale-mysql-database-backend-with-google-cloud-sql-and-proxysql#:~:text=Common%20approaches%20for%20horizontally%20scaling,with%20Cloud%20SQL%20and%20ProxySQL.The article only mentions

sharding as a concept, and not a solution for cloudsql.https://cloud.google.com/community/tutorials/horizontally-scale-mysql-database-backend-with-google-cloud-sql-and-proxysql#:~:text=SQL%20and%20ProxySQL.-,Sharding,logic%20or%20a%20query%20router. You can shard MySQL. Answer should be A.

upvoted 3 times

jay9114 3 years, 10 months ago

Ii Kk Kk You can shard cloudsql. Review this article - https://cloud.google.com/community/tutorials/horizontally-scale-mysql-database-backend-with-google-cloud-sql-and-proxysql#:~:text=Common%20approaches%20for%20horizontally%20scaling,with%20Cloud%20SQL%20and%20ProxySQL.

upvoted 4 times

jay9114 3 years, 10 months ago

Ii Kk Kk You can shard cloudsql. Review this article - https://cloud.google.com/community/tutorials/horizontally-scale-mysql-database-backend-with-google-cloud-sql-and-proxysql#:~:text=Common%20approaches%20for%20horizontally%20scaling,with%20Cloud%20SQL%20and%20ProxySQL.The article only mentions sharding as a concept, and not a solution for cloudsql.

upvoted 4 times

fiercedog 3 years, 8 months ago

The article only mentions sharding as a concept, and not a solution for cloudsql.

upvoted 1 times

Deb2293 3 years, 5 months ago

https://cloud.google.com/community/tutorials/horizontally-scale-mysql-database-backend-with-google-cloud-sql-and-proxysql#:~:text=SQL%20and%20ProxySQL.-,Sharding,logic%20or%20a%20query%20router. You can shard MySQL. Answer should be A.

upvoted 4 times

mj20201 4 years, 5 months ago

vote for A

upvoted 2 times

haroldbenites 4 years, 8 months ago

Go for A

upvoted 2 times

vincy2202 4 years, 9 months ago

A is the correct answer

upvoted 2 times

amxexam 4 years, 11 months ago

We can directly eliminate C and D we are doing some work that is already automated.

Still, I cannot make a point why not B is better than A? I believe adding memcash will give an additional boost

Can someone help me point out why A is better than B?

upvoted 4 times

amxexam 4 years, 11 months ago

We can directly eliminate C and D we are doing some work that is already automated. Still, I cannot make a point why not B is better than A? I believe adding memcash will give an additional boost Can someone help me point out why A is better than B?I suppose B is not a better option, since it indicates 'add 32core cpu', with no info of the current usage that seems like a over-kill.I would say only because of the below line| "You want to scale as usage increases" Line 1 Creating a 32 core machine upfront where we do not know what was the source machine cores would not be ideal . in that situation i would go with A

upvoted 4 times

cotam 4 years, 10 months ago

I suppose B is not a better option, since it indicates 'add 32core cpu', with no info of the current usage that seems like a over-kill.

upvoted 5 times

Ishu_awsguy 4 years ago

I would say only because of the below line| "You want to scale as usage increases" Line 1 Creating a 32 core machine upfront where we do not know what was the source machine cores would not be ideal . in that situation i would go with A

upvoted 4 times

MamthaSJ 5 years, 1 month ago

Answer is A

upvoted 4 times

Question #74 *Topic 1*

You are tasked with building an online analytical processing (OLAP) marketing analytics and reporting tool. This requires a relational database that can operate on hundreds of terabytes of data. What is the Google-recommended tool for such applications?

A. Cloud Spanner, because it is globally distributed

B. Cloud SQL, because it is a fully managed relational database

C. Cloud Firestore, because it offers real-time synchronization across devices

D. BigQuery, because it is designed for large-scale processing of tabular data

Suggested Answer: *D*

Reference: https://cloud.google.com/files/BigQueryTechnicalWP.pdf

Community vote distribution

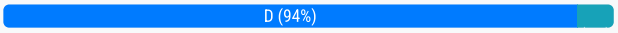

AWS56 Highly Voted 6 years, 7 months ago

Agree DD is okWhat about the relational part? BigQuery uses SQL but it's not relational... I'm not sure its DBigQuery is relational!Pls do not confuse - Cloud SQL and BigQuery are RDBMS. Cloud Datastore, Bigtable are NoSQL. Right answer is D - BQYes it is DD, OLAP=BQWell SaidThe question asks "This requires a relational database that can operate on hundreds of terabytes of data", but bq doesn't meet this condition?BigQuery supports relational and query of join tables.

upvoted 21 times

AWS56 Highly Voted 6 years, 7 months ago

Agree D

upvoted 21 times

tartar 6 years ago

D is ok

upvoted 5 times

tartar 6 years ago

D is okWhat about the relational part? BigQuery uses SQL but it's not relational... I'm not sure its DBigQuery is relational!Pls do not confuse - Cloud SQL and BigQuery are RDBMS. Cloud Datastore, Bigtable are NoSQL. Right answer is D - BQ

upvoted 5 times

Nastrand 5 years, 6 months ago

What about the relational part? BigQuery uses SQL but it's not relational... I'm not sure its D

upvoted 4 times

Nastrand 5 years, 6 months ago

What about the relational part? BigQuery uses SQL but it's not relational... I'm not sure its DBigQuery is relational!Pls do not confuse - Cloud SQL and BigQuery are RDBMS. Cloud Datastore, Bigtable are NoSQL. Right answer is D - BQ

upvoted 4 times

riflerrick 5 years, 2 months ago

BigQuery is relational!

upvoted 5 times

lovingsmart2000 5 years, 1 month ago

Pls do not confuse - Cloud SQL and BigQuery are RDBMS. Cloud Datastore, Bigtable are NoSQL. Right answer is D - BQ

upvoted 15 times

kumarp6 5 years, 9 months ago

Yes it is D

upvoted 2 times

nitinz 5 years, 5 months ago

D, OLAP=BQWell Said

upvoted 5 times

nitinz 5 years, 5 months ago

D, OLAP=BQ

upvoted 5 times

Sur_Nikki 3 years, 3 months ago

Well Said

upvoted 1 times

JasonL_GCP 4 years, 10 months ago

The question asks "This requires a relational database that can operate on hundreds of terabytes of data", but bq doesn't meet this condition?

upvoted 2 times

JasonL_GCP 4 years, 10 months ago

The question asks "This requires a relational database that can operate on hundreds of terabytes of data", but bq doesn't meet this condition?
BigQuery supports relational and query of join tables.

upvoted 2 times

elaineshi 4 years, 2 months ago

BigQuery supports relational and query of join tables.

upvoted 2 times

gfhbox0083 6 years, 2 months ago

D, for sure. BigQuery for OLAP Google Cloud Spanner for OLTP.

upvoted 18 times

gracjanborowiak Most Recent 2 years, 1 month ago

sql is oltp olap is data cube. lots of data which we try to process somehow. biq query is for that. firestore is for mobile and web apps. not so fast nosql db.

upvoted 1 times

Anandmrk 2 years, 5 months ago

Agreed

upvoted 1 times

AdityaGupta 2 years, 10 months ago

4 reasons to choose BQ (Supports Petabytes of data) - OLAP Data - Relational DB (SQL) - 100s of TB data - Analystics and Reporting

upvoted 1 times

Ashish1995 3 years, 5 months ago

D is obvious

upvoted 1 times

CGS22 3 years, 5 months ago

D is the right one

upvoted 1 times

SudhirAhirkar 3 years, 8 months ago

Cloud SQL/Spanner is OLTP DB but not OLAP. BQ is a well-known OLAP for analytics and also supports RBMS feature too... so I would got with D

upvoted 1 times

AniketD 3 years, 9 months ago

D is correct. BigQuery is relational. Cloud SQL is not OLAP; moreover it can not store/process hundreds of TB of data. Max size is 64 TB only.

upvoted 1 times

megumin 3 years, 9 months ago

D is ok

upvoted 1 times

SerGCP 3 years, 9 months ago

https://cloud.google.com/products/databases.

upvoted 1 times

zr79 3 years, 10 months ago

The words you need to focus "You are tasked with building an online analytical processing (OLAP) marketing analytics and reporting tool" which is BigQuery

upvoted 1 times

AzureDP900 3 years, 10 months ago

Big Query for large analytics , D is right

upvoted 1 times

Andre777 3 years, 11 months ago

The keyword in this context is OLAP. CloudSQL is Relational SQL for OLTP. Capacity wise, BQ supports for PB+ while CloudSQL only have max capacity of up to ~10TB. Again the questions specifically mention "hundreds of TB of data". So D is the answer.

upvoted 2 times

deepdowndave 3 years, 11 months ago

Why is it not CloudSQL? It supports TB data storage and the question is about a relational database, not a data warehouse such as BigQuery.The keyword in this context is OLAP. CloudSQL is Relational SQL for OLTP. Capacity wise, BQ supports for PB+ while CloudSQL only have max capacity of up to ~10TB. Again the questions specifically mention "hundreds of TB of data". So D is the answer.

upvoted 1 times

deepdowndave 3 years, 11 months ago

Why is it not CloudSQL? It supports TB data storage and the question is about a relational database, not a data warehouse such as BigQuery.

upvoted 1 times

alexandercamachop 3 years, 11 months ago

The answer is Big Query, D Secret: Analytical, Hundreds of TBTs. Relational. All of this are strictly meet by Big Query, if it had not said Analytical but rather, other keywords like High Availability then Cloud Spanner.

upvoted 1 times

Thornadoo 4 years ago

Guys, this is easy: OLTP - Cloud Spanner & Cloud SQL OLAP - Big Query NoSQL - Filestore and Big Table So answer is D.

upvoted 1 times

Question #75 *Topic 1*

You have deployed an application to Google Kubernetes Engine (GKE), and are using the Cloud SQL proxy container to make the Cloud SQL database available to the services running on Kubernetes. You are notified that the application is reporting database connection issues. Your company policies require a post- mortem. What should you do?

A. Use gcloud sql instances restart.

B. Validate that the Service Account used by the Cloud SQL proxy container still has the Cloud Build Editor role.

C. In the GCP Console, navigate to Stackdriver Logging. Consult logs for (GKE) and Cloud SQL.

D. In the GCP Console, navigate to Cloud SQL. Restore the latest backup. Use kubectl to restart all pods.

Suggested Answer: *C*
Official Answer: C

Community vote distribution
C (100%)

jcmoranp Highly Voted 5 years, 10 months ago
post mortem always includes log analysis, answer is CAGREE CC is right for Root Cause Analysis.Thanks for the info
upvoted 68 times

jcmoranp Highly Voted 5 years, 10 months ago
post mortem always includes log analysis, answer is C
upvoted 68 times

AWS56 5 years, 7 months ago
AGREE C
upvoted 3 times

AzureDP900 2 years, 10 months ago
C is right for Root Cause Analysis.
upvoted 1 times

Sur_Nikki 2 years, 3 months ago
Thanks for the info
upvoted 1 times

MamthaSJ 4 years, 1 month ago
Answer is C
upvoted 5 times

pakilodi Most Recent 1 year, 8 months ago
C -> post-mortem = log analysis
upvoted 2 times

AdityaGupta 1 year, 10 months ago
You can jump on to the conlusion hence answer is not B. Consulting logs is always a good way to start investigation. and A and D is not a choice.
upvoted 1 times

ale_brd_111 2 years, 9 months ago
Stackdriver is deprecated, now you must navigate to Cloud Logging.
upvoted 2 times

megumin 2 years, 9 months ago
C is ok
upvoted 1 times

Mahmoud_E 2 years, 10 months ago
C is the right answer
upvoted 1 times

Jay_Krish 2 years, 11 months ago
Logical answer is C. But is Stackdriver Logging enabled by default? Appreciate if someone could answer this?
upvoted 1 times

haroldbenites 3 years, 8 months ago

Go for C

upvoted 1 times

pakilodi 3 years, 8 months ago

post mortem = logs

upvoted 1 times

vincy2202 3 years, 8 months ago

C is the correct answer

upvoted 1 times

joe2211 3 years, 8 months ago

vote C

upvoted 1 times

MaxNRG 3 years, 9 months ago

C – in GCP Console navigate to Stackdriver Logging. Consult logs for Kubernetes Engine and Cloud SQL. A/D – is an immediate attempt to fix an issue. No analysis. B – is irrelevant at all. Cloud SQL proxy should not build anything in production.

upvoted 4 times

lovingsmart2000 4 years, 1 month ago

Answer is C. I request all here - not to blindly follow the answers published at coursera or udemy as most of them are copy-pasted answer and are not real. Examtopis provides the more accurate answers and also support with comments

upvoted 4 times

lovingsmart2000 4 years, 1 month ago

Answer is B. I request all here - not to blindly follow the answers published at coursera or udemy as most of them are copy-pasted answer and are not real. Examtopis provides the more accurate answers and also support with comments

upvoted 2 times

lovingsmart2000 4 years, 1 month ago

Answer is B. I request all here - not to blindly follow the answers published at coursera or udemy as most of them are copy-pasted answer and are not real. Examtopis provides the more accurate answers and also support with commentsWhy Service Account needs Cloud Build Editor role for accessing Cloud SQL? The role is misleading/wrong, so B is wrong.

upvoted 2 times

ashish_t 3 years, 10 months ago

Why Service Account needs Cloud Build Editor role for accessing Cloud SQL? The role is misleading/wrong, so B is wrong.

upvoted 4 times

victory108 4 years, 3 months ago

C. In the GCP Console, navigate to Stackdriver Logging. Consult logs for Kubernetes Engine and Cloud SQL.

upvoted 2 times

un 4 years, 3 months ago

C is correct

upvoted 1 times

Question #76 *Topic 1*

Your company pushes batches of sensitive transaction data from its application server VMs to Cloud Pub/Sub for processing and storage. What is the Google- recommended way for your application to authenticate to the required Google Cloud services?

A. Ensure that VM service accounts are granted the appropriate Cloud Pub/Sub IAM roles.

B. Ensure that VM service accounts do not have access to Cloud Pub/Sub, and use VM access scopes to grant the appropriate Cloud Pub/Sub IAM roles.

C. Generate an OAuth2 access token for accessing Cloud Pub/Sub, encrypt it, and store it in Cloud Storage for access from each VM.

D. Create a gateway to Cloud Pub/Sub using a Cloud Function, and grant the Cloud Function service account the appropriate Cloud Pub/Sub IAM roles.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (100%)

AWS56 Highly Voted 6 years, 1 month ago
Agree AYes A it isA is correct
upvoted 26 times

AWS56 Highly Voted 6 years, 1 month ago
Agree A
upvoted 26 times

kumarp6 5 years, 3 months ago
Yes A it is
upvoted 2 times

nitinz 4 years, 11 months ago
A is correct
upvoted 2 times

JustJack21 4 years, 5 months ago
It's because of questions like these that I do not feel guilty about using question banks :D In what world would you accept value requirements like this from your user? Wouldn't you ask "Do you want to just authenticate? or the data to be encrypted on its way to pub/sub?" I'll ignore the first part of the question and assume all data is sensitive, and focus on "What is the Google- recommended way for your application to authenticate to the required Google Cloud services?" -- The answer then is A.

Use encryption and defense-in-depth for the first part.
upvoted 13 times

JustJack21 4 years, 5 months ago
It's because of questions like these that I do not feel guilty about using question banks :D In what world would you accept value requirements like this from your user? Wouldn't you ask "Do you want to just authenticate? or the data to be encrypted on its way to pub/sub?" I'll ignore the first part of the question and assume all data is sensitive, and focus on "What is the Google- recommended way for your application to authenticate to the required Google Cloud services?" -- The answer then is A. Use encryption and defense-in-depth for the first part.Service accounts use keys> It's because of questions like these that I do not feel guilty about using question banks :D Same. To me, it wasn't clear whether the servers were in google or not due to the question about accessing google cloud. It was asked as if the VMs were outside of google
upvoted 13 times

AMEJack 3 years, 4 months ago
Service accounts use keys
upvoted 1 times

bandegg Most Recent 2 years, 1 month ago
> It's because of questions like these that I do not feel guilty about using question banks :D Same. To me, it wasn't clear whether the servers were in google or not due to the question about accessing google cloud. It was asked as if the VMs were outside of google
upvoted 3 times

red_panda 2 years, 8 months ago

A is correct for me. It's batch, so no cloud function

upvoted 2 times

omermahgoub 3 years, 1 month ago

A. Ensure that VM service accounts are granted the appropriate Cloud Pub/Sub IAM roles.

The Google-recommended way for your application to authenticate to Cloud Pub/Sub and other Google Cloud services when running on Compute Engine VMs is to use VM service accounts. VM service accounts are automatically created when you create a Compute Engine VM, and they are associated with the VM instance. To authenticate to Cloud Pub/Sub and other Google Cloud services, you should ensure that the VM service accounts are granted the appropriate IAM roles.

upvoted 9 times

omermahgoub 3 years, 1 month ago

A. Ensure that VM service accounts are granted the appropriate Cloud Pub/Sub IAM roles. The Google-recommended way for your application to authenticate to Cloud Pub/Sub and other Google Cloud services when running on Compute Engine VMs is to use VM service accounts. VM service accounts are automatically created when you create a Compute Engine VM, and they are associated with the VM instance. To authenticate to Cloud Pub/Sub and other Google Cloud services, you should ensure that the VM service accounts are granted the appropriate IAM roles.Option B, ensuring that VM service accounts do not have access to Cloud Pub/Sub and using VM access scopes to grant the appropriate Cloud Pub/Sub IAM roles, would not be a suitable solution because VM service accounts are required for authentication to Google Cloud services. Option C, generating an OAuth2 access token for accessing Cloud Pub/Sub, encrypting it, and storing it in Cloud Storage for access from each VM, would not be a suitable solution because it would require manual management of access tokens, which can be error-prone and insecure. Option D, creating a gateway to Cloud Pub/Sub using a Cloud Function and granting the Cloud Function service account the appropriate Cloud Pub/Sub IAM roles, would not be a suitable solution because it would not allow the application to directly authenticate to Cloud Pub/Sub.Great way of explanation..By removing/elimination approach

upvoted 9 times

omermahgoub 3 years, 1 month ago

Option B, ensuring that VM service accounts do not have access to Cloud Pub/Sub and using VM access scopes to grant the appropriate Cloud Pub/Sub IAM roles, would not be a suitable solution because VM service accounts are required for authentication to Google Cloud services. Option C, generating an OAuth2 access token for accessing Cloud Pub/Sub, encrypting it, and storing it in Cloud Storage for access from each VM, would not be a suitable solution because it would require manual management of access tokens, which can be error-prone and insecure. Option D, creating a gateway to Cloud Pub/Sub using a Cloud Function and granting the Cloud Function service account the appropriate Cloud Pub/Sub IAM roles, would not be a suitable solution because it would not allow the application to directly authenticate to Cloud Pub/Sub.Great way of explanation..By removing/elimination approach

upvoted 4 times

omermahgoub 3 years, 1 month ago

Option B, ensuring that VM service accounts do not have access to Cloud Pub/Sub and using VM access scopes to grant the appropriate Cloud Pub/Sub IAM roles, would not be a suitable solution because VM service accounts are required for authentication to Google Cloud services.

Option C, generating an OAuth2 access token for accessing Cloud Pub/Sub, encrypting it, and storing it in Cloud Storage for access from each VM, would not be a suitable solution because it would require manual management of access tokens, which can be error-prone and insecure.

Option D, creating a gateway to Cloud Pub/Sub using a Cloud Function and granting the Cloud Function service account the appropriate Cloud Pub/Sub IAM roles, would not be a suitable solution because it would not allow the application to directly authenticate to Cloud Pub/Sub.

upvoted 4 times

Sur_Nikki 2 years, 9 months ago

Great way of explanation..By removing/elimination approach

upvoted 1 times

megumin 3 years, 3 months ago

A is ok

upvoted 1 times

Mahmoud_E 3 years, 3 months ago

A is the correct answer

upvoted 1 times

DrishaS4 3 years, 6 months ago

https://cloud.google.com/iam/docs/understanding-service-accounts

upvoted 1 times

Pazzooo 4 years ago

The combination of Roles assigned to Service accounts granted to VMs is the way to go. :)

upvoted 2 times

elenamatay 4 years, 1 month ago

Service accounts are recommended for almost all cases in Pub/Sub (see https://cloud.google.com/pubsub/docs/authentication#service-accounts)

upvoted 3 times

haroldbenites 4 years, 2 months ago

Go for A.

upvoted 2 times

vincy2202 4 years, 2 months ago

A is the correct answer

upvoted 1 times

MaxNRG 4 years, 3 months ago

A – ensure that VM service accounts are granted the appropriate Cloud Pub/Sub IAM roles. Check Migrating Data to GCP section of this page: https://cloud.google.com/iam/docs/understanding-service-accounts You will create a service account key and use it from an external process to call Cloud Platform APIs.

upvoted 3 times

Bakili 4 years, 3 months ago

A is very correct

upvoted 1 times

MamthaSJ 4 years, 7 months ago

Answer is A

upvoted 2 times

victory108 4 years, 9 months ago

A. Ensure that VM service accounts are granted the appropriate Cloud Pub/Sub IAM roles.Agreed with A

upvoted 3 times

victory108 4 years, 9 months ago

A. Ensure that VM service accounts are granted the appropriate Cloud Pub/Sub IAM roles.

upvoted 3 times

AzureDP900 3 years, 4 months ago

Agreed with A

upvoted 1 times

un 4 years, 9 months ago

A is correct

upvoted 1 times

kartikjena31 4 years, 10 months ago

Ans. A

upvoted 1 times

Question #77 *Topic 1*

You want to establish a Compute Engine application in a single VPC across two regions. The application must communicate over VPN to an on-premises network.
How should you deploy the VPN?

A. Use VPC Network Peering between the VPC and the on-premises network.

B. Expose the VPC to the on-premises network using IAM and VPC Sharing.

C. Create a global Cloud VPN Gateway with VPN tunnels from each region to the on-premises peer gateway.

D. Deploy Cloud VPN Gateway in each region. Ensure that each region has at least one VPN tunnel to the on-premises peer gateway.

Suggested Answer: *D*
Official Answer: D

Community vote distribution

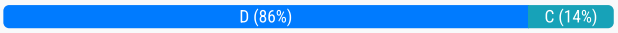

Googler2 Highly Voted 6 years, 4 months ago
It can't be -A - VPC Network Peering only allows private RFC 1918 connectivity across two Virtual Private Cloud (VPC) networks. In this example is one VPC with on-premise network https://cloud.google.com/vpc/docs/vpc-peering

It is not definitely - B - Can't be

It is not C - Because Cloud VPN gateways and tunnels are regional objects, not global

So, it the answer is D - https://cloud.google.com/vpn/docs/how-to/creating-static-vpns
upvoted 48 times

Googler2 Highly Voted 6 years, 4 months ago
It can't be -A - VPC Network Peering only allows private RFC 1918 connectivity across two Virtual Private Cloud (VPC) networks. In this example is one VPC with on-premise network https://cloud.google.com/vpc/docs/vpc-peering It is not definitely - B - Can't be It is not C - Because Cloud VPN gateways and tunnels are regional objects, not global So, it the answer is D - https://cloud.google.com/vpn/docs/how-to/creating-static-vpnsWhy not A? https://cloud.google.com/vpc/docs/vpc-peering#benefits_of_exchanging_custom_routes The second use case is exactly what is in the question. Don't get the argument about RFC 1918. Will go with Ahttps://cloud.google.com/vpc/docs/vpc-peering allows internal IP address connectivity across two VPC so A is not the answer as the on premise network need to use public IP. cmiiwThe question clearly asks us to use VPN.Agreed with D.
upvoted 48 times

amxexam 4 years, 11 months ago
Why not A? https://cloud.google.com/vpc/docs/vpc-peering#benefits_of_exchanging_custom_routes The second use case is exactly what is in the question.

Don't get the argument about RFC 1918. Will go with A
upvoted 1 times

amxexam 4 years, 11 months ago
Why not A? https://cloud.google.com/vpc/docs/vpc-peering#benefits_of_exchanging_custom_routes The second use case is exactly what is in the question. Don't get the argument about RFC 1918. Will go with Ahttps://cloud.google.com/vpc/docs/vpc-peering allows internal IP address connectivity across two VPC so A is not the answer as the on premise network need to use public IP. cmiiwThe question clearly asks us to use VPN.
upvoted 1 times

ochanz 4 years, 8 months ago
https://cloud.google.com/vpc/docs/vpc-peering allows internal IP address connectivity across two VPC so A is not the answer as the on premise network need to use public IP. cmiiw
upvoted 4 times

AdityaGupta 2 years, 10 months ago
The question clearly asks us to use VPN.
upvoted 2 times

AzureDP900 3 years, 10 months ago

Agreed with D.

upvoted 1 times

TaherShaker 5 years, 9 months ago

Just Passed my exam and I answered (D) for this questionsound promising dudeIS the Exam Idea questions enough dude, for passing this exam?

upvoted 21 times

TaherShaker 5 years, 9 months ago

Just Passed my exam and I answered (D) for this question

upvoted 21 times

M_Asep 4 years, 8 months ago

sound promising dude

upvoted 3 times

Sur_Nikki 3 years, 3 months ago

IS the Exam Idea questions enough dude, for passing this exam?

upvoted 3 times

Snafkin547 Most Recent 1 year, 2 months ago

You don't need to deploy VPN to two regions, because of its global nature. C is correct

upvoted 1 times

ccpmad 2 years, 1 month ago

Option C: Create a global VPN gateway and establish VPN tunnels from each region to the on-premises peer gateway. This suggests that a single global VPN gateway manages the tunnels from both regions. Option D: Deploy a VPN gateway in each region and ensure that each region has at least one VPN tunnel to the on-premises peer gateway. This indicates that each region has its own VPN gateway. >Option D ensures that there is a VPN gateway in each region, providing greater redundancy. If a gateway in one region fails, the gateway in the other region remains operational.

upvoted 2 times

santoshchauhan 2 years, 4 months ago

Global Cloud VPN Gateway: This feature allows for the creation of a single VPN gateway that can serve multiple regions within the same VPC network. By creating a global VPN gateway, you can efficiently manage VPN connections from all regions of your VPC to your on-premises network. Simplicity and Efficiency: Using a global gateway simplifies the configuration and management of VPN connections as opposed to maintaining separate regional VPN gateways. It centralizes the VPN endpoint on the Google Cloud side, reducing the complexity of the network setup. Reliable and Secure Communication: The global Cloud VPN Gateway allows for secure, encrypted tunnels between Google Cloud and the on-premises network, ensuring that the application's inter-regional and on-premises communications are secure.

upvoted 2 times

salvo007 2 years, 7 months ago

C is wrong. A global vpn is a single region resource. https://cloud.google.com/network-connectivity/docs/vpn/how-to/creating-ha-vpn?hl=it gcloud compute vpn-gateways create GW_NAME \ --network=NETWORK \ --region=REGION \ --stack-type=IP_STACK so D is the answer

upvoted 2 times

gcmrjbr 2 years, 7 months ago

It´s option C! So, while the VPN Gateway itself is a regional resource, its scope can be effectively global as it can serve resources across different regions within the same Virtual Private Cloud (VPC). This is why it's sometimes referred to as a 'global' service in the context of its functionality, even though strictly speaking, it's a regional resource.

upvoted 2 times

AdityaGupta 2 years, 10 months ago

Each Cloud VPN gateway is a regional resource that uses one or more regional external IP addresses. A Cloud VPN gateway can connect to a peer VPN gateway.

upvoted 2 times

LaxmanTiwari 3 years, 3 months ago

It can't be -A - VPC Network Peering only allows private RFC 1918 connectivity across two Virtual Private Cloud (VPC) networks. In this example is one VPC with on-premise network https://cloud.google.com/vpc/docs/vpc-peering It is not definitely - B - Can't be It is not C - Because Cloud VPN gateways and tunnels are regional objects, not global So, it the answer is D - https://cloud.google.com/vpn/docs/how-to/creating-static-vpn

upvoted 3 times

vvkds 3 years, 7 months ago

D looks fine.

upvoted 1 times

oms_muc 3 years, 8 months ago

As HA isn't required, why do we need two VPN gateways?

upvoted 2 times

megumin 3 years, 9 months ago

D is ok

upvoted 1 times

Mahmoud_E 3 years, 10 months ago

D is the correct answer, in order to do A you will need VPN., or interconnect

upvoted 1 times

zr79 3 years, 10 months ago

there is two VPN: 1. classic VPN 2. HA VPN

upvoted 1 times

DrishaS4 4 years ago

Cloud VPN Gateway is a regional service, not global.

upvoted 4 times

elaineshi 4 years, 2 months ago

Why not C? services across regions can communicate to each other, VPN only connects to the closet region, and all the VPC shall be connected if firewall's set.

upvoted 2 times

haroldbenites 4 years, 8 months ago

Go for D. Cloud VPN Gateway is regional. NOt Global gcloud compute vpn-gateways create GW_NAME \ --network=NETWORK \ --region=REGION

upvoted 2 times

Question #78 *Topic 1*

Your applications will be writing their logs to BigQuery for analysis. Each application should have its own table. Any logs older than 45 days should be removed.
You want to optimize storage and follow Google-recommended practices. What should you do?

A. Configure the expiration time for your tables at 45 days

B. Make the tables time-partitioned, and configure the partition expiration at 45 days

C. Rely on BigQuery's default behavior to prune application logs older than 45 days

D. Create a script that uses the BigQuery command line tool (bq) to remove records older than 45 days

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (100%)

KouShikyou Highly Voted 6 years, 4 months ago
Could you please help clarify? I think B is correct. It looks like table will be deleted with option A. https://cloud.google.com/bigquery/docs/managing-tables#updating_a_tables_expiration_time When you delete a table, any data in the table is also deleted. To automatically delete tables after a specified period of time, set the default table expiration for the dataset or set the expiration time when you create the table.
upvoted 40 times

KouShikyou Highly Voted 6 years, 4 months ago
Could you please help clarify? I think B is correct. It looks like table will be deleted with option A. https://cloud.google.com/bigquery/docs/managing-tables#updating_a_tables_expiration_time When you delete a table, any data in the table is also deleted. To automatically delete tables after a specified period of time, set the default table expiration for the dataset or set the expiration time when you create the table.B is okB partition tableit is B, if you use option A, on 46th day there is no table/content in table for application :)Agreed and going with B
upvoted 40 times

tartar 5 years, 6 months ago
B is ok
upvoted 8 times

nitinz 4 years, 11 months ago
B partition table
upvoted 4 times

kumarp6 5 years, 3 months ago
it is B, if you use option A, on 46th day there is no table/content in table for application :)
upvoted 11 times

AzureDP900 3 years, 4 months ago
Agreed and going with B
upvoted 2 times

aviv 6 years, 2 months ago
Agreed with B.
upvoted 10 times

OSAMA911 Most Recent 1 year, 12 months ago
I think B is correct.
upvoted 1 times

AdityaGupta 2 years, 4 months ago
https://cloud.google.com/bigquery/docs/managing-partitioned-tables#partition-expiration B is the correct answer.
upvoted 4 times

SSPPJi 2 years, 7 months ago
https://cloud.google.com/bigquery/docs/managing-partitioned-tables#partition-expiration
upvoted 4 times

FaizAhmed 2 years, 7 months ago

B is correct

upvoted 1 times

Sur_Nikki 2 years, 9 months ago

B seems correct as this will partitioning will create a filter criteria on the basis of which specified actions on logs will be taken

upvoted 1 times

examch 3 years, 1 month ago

B is the correct answer, If your tables are partitioned by date, the dataset's default table expiration applies to the individual partitions. You can also control partition expiration using the time_partitioning_expiration flag in the bq command-line tool or the expirationMs configuration setting in the API. When a partition expires, data in the partition is deleted but the partitioned table is not dropped even if the table is empty.
https://cloud.google.com/bigquery/docs/best-practices-storage

upvoted 7 times

megumin 3 years, 3 months ago

B is ok

upvoted 1 times

MarcoEscanor 3 years, 3 months ago

B - You can control partition expiration using the time_partitioning_expiration flag in the bq command-line
https://cloud.google.com/bigquery/docs/best-practices-storage

upvoted 2 times

AhmedH7793 3 years, 5 months ago

B is okay

upvoted 1 times

DrishaS4 3 years, 6 months ago

Using Table-Partitions.

upvoted 1 times

haroldbenites 4 years, 2 months ago

Go for B. https://cloud.google.com/bigquery/docs/creating-partitioned-tables#sql CREATE TABLE mydataset.newtable (transaction_id INT64, transaction_date DATE) PARTITION BY transaction_date OPTIONS(partition_expiration_days=3, require_partition_filter=true)

upvoted 2 times

vincy2202 4 years, 2 months ago

B is the correct answer

upvoted 2 times

MaxNRG 4 years, 3 months ago

B – Make the tables time-partitioned and configure the partition expiration at 45 days. A – if you use table expiration time, then it will remove the whole table after 45 days. D – requires extra work and is not automatic.

upvoted 2 times

Unfaithful 4 years, 6 months ago

Answer: B Support: https://cloud.google.com/bigquery/docs/best-practices-storage

upvoted 1 times

Question #79 *Topic 1*

You want your Google Kubernetes Engine cluster to automatically add or remove nodes based on CPU load.
What should you do?

A. Configure a HorizontalPodAutoscaler with a target CPU usage. Enable the Cluster Autoscaler from the GCP Console.

B. Configure a HorizontalPodAutoscaler with a target CPU usage. Enable autoscaling on the managed instance group for the cluster using the gcloud command.

C. Create a deployment and set the maxUnavailable and maxSurge properties. Enable the Cluster Autoscaler using the gcloud command.

D. Create a deployment and set the maxUnavailable and maxSurge properties. Enable autoscaling on the cluster managed instance group from the GCP Console.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (100%)

Unfaithful Highly Voted 4 years ago
Answer: A Support: How does Horizontal Pod Autoscaler work with Cluster Autoscaler?

Horizontal Pod Autoscaler changes the deployment's or replicaset's number of replicas based on the current CPU load. If the load increases, HPA will create new replicas, for which there may or may not be enough space in the cluster. If there are not enough resources, CA will try to bring up some nodes, so that the HPA-created pods have a place to run. If the load decreases, HPA will stop some of the replicas. As a result, some nodes may become underutilized or completely empty, and then CA will terminate such unneeded nodes.
upvoted 65 times

Unfaithful Highly Voted 4 years ago
Answer: A Support: How does Horizontal Pod Autoscaler work with Cluster Autoscaler? Horizontal Pod Autoscaler changes the deployment's or replicaset's number of replicas based on the current CPU load. If the load increases, HPA will create new replicas, for which there may or may not be enough space in the cluster. If there are not enough resources, CA will try to bring up some nodes, so that the HPA-created pods have a place to run. If the load decreases, HPA will stop some of the replicas. As a result, some nodes may become underutilized or completely empty, and then CA will terminate such unneeded nodes.Good ExplainationNice and detailed explanation. I agree with A.Nice and detailed explanation. I agree with A.very well explained
upvoted 65 times

Rajasa 3 years, 8 months ago
Good Explaination
upvoted 3 times

AzureDP900 2 years, 10 months ago
Nice and detailed explanation. I agree with A.
upvoted 1 times

heretolearnazure Most Recent 1 year, 11 months ago
very well explained
upvoted 1 times

natpilot 5 years, 6 months ago
i'm for A, but the question in ambiguous, because requires the autoscale of nodes (not pod) when the cpu overload, but in answer use k8s pod autoscaler based on cpu load (cpu load for pod, not nodes). strange
upvoted 25 times

natpilot 5 years, 6 months ago
i'm for A, but the question in ambiguous, because requires the autoscale of nodes (not pod) when the cpu overload, but in answer use k8s pod autoscaler based on cpu load (cpu load for pod, not nodes). strangeConfuse with the question like you mentioned. Autoscale is via nodes not pod.. and can only be configure using gcloud command.Agreed, the question is not about pods, but answers are also talking about pods (not only) A is correct because B is wrong according to https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler "Caution: Do not enable

Compute Engine autoscaling for managed instance groups for your cluster nodes. GKE's cluster autoscaler is separate from Compute Engine autoscaling"

upvoted 25 times

skywalker 5 years, 3 months ago

Confuse with the question like you mentioned. Autoscale is via nodes not pod.. and can only be configure using gcloud command.

upvoted 6 times

p4 4 years, 9 months ago

Agreed, the question is not about pods, but answers are also talking about pods (not only) A is correct because B is wrong according to https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler "Caution: Do not enable Compute Engine autoscaling for managed instance groups for your cluster nodes. GKE's cluster autoscaler is separate from Compute Engine autoscaling"

upvoted 21 times

LaxmanTiwari 2 years, 3 months ago

Nice and detailed explanation. I agree with A.

upvoted 1 times

Sur_Nikki 2 years, 3 months ago

A seems correct. y to create managed instance groups unnecessarily?

upvoted 1 times

Deb2293 2 years, 5 months ago

The answer is A. More nodes mean it's horizontal scaling (increase VMs means vertical scaling of infrastructure). Cluster AutoScalar is used for increasing number of nodes.

upvoted 1 times

examch 2 years, 7 months ago

A is the Correct answer, Horizontal Pod Autoscaler and Cluster Autoscaler can be used together to provision new pods and new nodes as per the CPU utilization. https://www.youtube.com/watch?v=VNAWA6NkoBs

upvoted 2 times

megumin 2 years, 9 months ago

ok for A

upvoted 1 times

Rajeev26 2 years, 10 months ago

MIG not for GKE as option B and C, D are not relevant to question

upvoted 1 times

abirroy 2 years, 11 months ago

Configure a HorizontalPodAutoscaler with a target CPU usage. Enable the Cluster Autoscaler from the GCP Console.

upvoted 1 times

gee1979 2 years, 11 months ago

A... The HPA and CA complement each other for truly efficient scaling. If the load increases, HPA will create new replicas. If there isn't enough space for these replicas, CA will provision some nodes, so that the HPA-created pods have a place to run. 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, or in response to custom metrics reported from within Kubernetes or external metrics from sources outside of your cluster.

upvoted 1 times

6721sora 2 years, 11 months ago

A is wrong. Pod scaling only spins up additional pods. Not nodes. Cluster Autoscaler does adding of nodes automatically. I am surprised that so many people think that A is the correct answer. Correct answer per me is C

upvoted 2 times

DrishaS4 3 years ago

Horizontal Pod Autoscaler changes the deployment's or replicaset's number of replicas based on the current CPU load. If the load increases, HPA will create new replicas, for which there may or may not be enough space in the cluster. If there are not enough resources, CA will try to bring up some nodes, so that the HPA-created pods have a place to run. If the load decreases, HPA will stop some of the replicas. As a result, some nodes may become underutilized or completely empty, and then CA will terminate such unneeded nodes.

upvoted 2 times

OrangeTiger 3 years, 7 months ago

I agree A is correct. I found quicklab. Understanding and Combining GKE Autoscaling Strategies.

upvoted 1 times

ehgm 3 years, 7 months ago

B and D: You must never change the GKE managed instance group. C and D: maxUnavailable and maxSurge are used for rolling update A. It is the correct.

upvoted 6 times

haroldbenites 3 years, 8 months ago

Go for A

upvoted 1 times

MaxNRG 3 years, 9 months ago

Create Horizontal Autoscaler (min,max for pods): kubectl autoscale deployment my-app --max 6 --min 4 --cpu-percent 50 Autoscaling cluster: gcloud container clusters create example-cluster \ --zone us-central1-a \ --node-locations us-central1-a,us-central1-b,us-central1-f \ --num-nodes 2 --enable-autoscaling --min-nodes 1 --max-nodes 4 Check scaling an application and Horizontal Pod Autoscaler: https://kubernetes.io/docs/tasks/run-application/horizontal-pod-autoscale/ Manual Cluster Resizing: https://cloud.google.com/kubernetes-engine/docs/how-to/resizing-a-cluster https://cloud.google.com/kubernetes-engine/docs/how-to/scaling-appsCorrect answer A.

upvoted 6 times

MaxNRG 3 years, 9 months ago

Create Horizontal Autoscaler (min,max for pods): kubectl autoscale deployment my-app --max 6 --min 4 --cpu-percent 50 Autoscaling cluster: gcloud container clusters create example-cluster \ --zone us-central1-a \ --node-locations us-central1-a,us-central1-b,us-central1-f \ --num-nodes 2 --enable-autoscaling --min-nodes 1 --max-nodes 4 Check scaling an application and Horizontal Pod Autoscaler: https://kubernetes.io/docs/tasks/run-application/horizontal-pod-autoscale/ Manual Cluster Resizing: https://cloud.google.com/kubernetes-engine/docs/how-to/resizing-a-cluster https://cloud.google.com/kubernetes-engine/docs/how-to/scaling-apps

upvoted 6 times

MaxNRG 3 years, 9 months ago

Correct answer A.

upvoted 1 times

MaxNRG 3 years, 9 months ago

D. Cloud VPN provides secure IPSec connection, though Direct Peering doesn't. Also, check selection diagram "What GCP connection is right for you?" on Hybrid Connectivity page. https://cloud.google.com/hybrid-connectivity/ It explicitly points that Cloud VPN and Dedicated Interconnect are for extension of you Data Center to Cloud (== of private compute resources). And Direct Peering for accessing GSuite (full set of GCP resources). Direct Peering: https://cloud.google.com/network-connectivity/docs/direct-peering Cloud VPN: https://cloud.google.com/network-connectivity/docs/vpn/concepts/overview Choose Inteconnect Type: https://cloud.google.com/network-connectivity/docs/how-to/choose-product#cloud-interconnect only suggests Dedicted/Partner and Cloud VPN. This Disaster Recovery scenario is described here, in section "Transferring data to and from GCP": https://cloud.google.com/architecture/dr-scenarios-building-blocks#transferring_data_to_and_from

upvoted 1 times

Question #80 *Topic 1*

You need to develop procedures to verify resilience of disaster recovery for remote recovery using GCP. Your production environment is hosted on-premises. You need to establish a secure, redundant connection between your on-premises network and the GCP network.
What should you do?

A. Verify that Dedicated Interconnect can replicate files to GCP. Verify that direct peering can establish a secure connection between your networks if Dedicated Interconnect fails.

B. Verify that Dedicated Interconnect can replicate files to GCP. Verify that Cloud VPN can establish a secure connection between your networks if Dedicated Interconnect fails.

C. Verify that the Transfer Appliance can replicate files to GCP. Verify that direct peering can establish a secure connection between your networks if the Transfer Appliance fails.

D. Verify that the Transfer Appliance can replicate files to GCP. Verify that Cloud VPN can establish a secure connection between your networks if the Transfer Appliance fails.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (87%) A (13%)

KouShikyou Highly Voted 6 years, 10 months ago
I think B is correct answer.
upvoted 44 times

KouShikyou Highly Voted 6 years, 10 months ago
I think B is correct answer.B is okIts quite a fun to use Transfer Appliance for DR, I think answer is BActually, how ca this be given as a option even?
only B works
upvoted 44 times

tartar 6 years ago
B is ok
upvoted 8 times

kumarp6 5 years, 9 months ago
Its quite a fun to use Transfer Appliance for DR, I think answer is BActually, how ca this be given as a option even?
upvoted 6 times

kumarp6 5 years, 9 months ago
Its quite a fun to use Transfer Appliance for DR, I think answer is B
upvoted 6 times

Sur_Nikki 3 years, 3 months ago
Actually, how ca this be given as a option even?
upvoted 1 times

nitinz 5 years, 5 months ago
only B works
upvoted 1 times

MeasService 6 years, 10 months ago
Agree B is correct. Transfer appliance is a physical appliance for transferring huge bulk of data. does not fit into disaster recovery testing. out of A and B, B seems to be more nearest answer. One would not have direct peering and Dedicated interconnect in a solution
upvoted 27 times

Ekramy_Elnaggar Most Recent 1 year, 9 months ago
Dedicated Interconnect as Primary , and Cloud VPN as Backup
upvoted 2 times

nareshthumma 1 year, 9 months ago
answer is B
upvoted 1 times

hitmax87 2 years, 3 months ago

I go to B, because direct peering anyway requires VPN connection if you want to get access to VPC.

upvoted 4 times

MahAli 2 years, 8 months ago

Selected Answer: A

Why you guys are choosing VPN? the reason to use Dedicated Interconnect is to have the max bandwidth available, does VPN give you that option in the first place? why not thinking about separate direct peering connection which might give a better performance than VPN?

upvoted 2 times

MahAli 2 years, 8 months ago

Why you guys are choosing VPN? the reason to use Dedicated Interconnect is to have the max bandwidth available, does VPN give you that option in the first place? why not thinking about separate direct peering connection which might give a better performance than VPN?BTW with direct peering you are going through the service provider network which makes more sense to get different connectivity optionDirect peering allows only on premises to connect to Google services in GCP . If needed to connect with Google workspace where all projects hosted they dedicated or partner interconnect is required. https://cloud.google.com/network-connectivity/docs/direct-peering B is the best answer

upvoted 2 times

MahAli 2 years, 8 months ago

BTW with direct peering you are going through the service provider network which makes more sense to get different connectivity option

upvoted 1 times

MahAli 2 years, 8 months ago

BTW with direct peering you are going through the service provider network which makes more sense to get different connectivity optionDirect peering allows only on premises to connect to Google services in GCP . If needed to connect with Google workspace where all projects hosted they dedicated or partner interconnect is required. https://cloud.google.com/network-connectivity/docs/direct-peering B is the best answer

upvoted 1 times

Diwz 2 years, 4 months ago

Direct peering allows only on premises to connect to Google services in GCP . If needed to connect with Google workspace where all projects hosted they dedicated or partner interconnect is required. https://cloud.google.com/network-connectivity/docs/direct-peering B is the best answer

upvoted 3 times

AdityaGupta 2 years, 10 months ago

You need to develop procedures to verify resilience of disaster recovery for remote recovery using GCP. Your production environment is hosted on-premises. You need to establish a secure, redundant connection between your on-premises network and the GCP network. What should you do? A. Verify that Dedicated Interconnect can replicate files to GCP. Verify that direct peering can establish a secure connection between your networks if Dedicated Interconnect fails. B. Verify that Dedicated Interconnect can replicate files to GCP. Verify that Cloud VPN can establish a secure connection between your networks if Dedicated Interconnect fails. Why Not A, as question asks "to establish a secure, redundant connection between your on-premises network and the GCP network." Is VPN considered more reliable than Direct Peering?? Both VPN and Direct Peering will provide redundant connection. I am not concerned about cost Direct Interconnect is already there.

upvoted 1 times

FaizAhmed 3 years, 1 month ago

B is right,

upvoted 1 times

omermahgoub 3 years, 8 months ago

The correct answer is B. Verify that Dedicated Interconnect can replicate files to GCP. Verify that Cloud VPN can establish a secure connection between your networks if Dedicated Interconnect fails.

Dedicated Interconnect is a connection that provides a private, dedicated connection between your on-premises network and GCP over a Google-owned network. It is a secure and reliable option for connecting your on-premises network to GCP. You can use it to replicate files to GCP as a part of your disaster recovery plan.

If Dedicated Interconnect fails for any reason, it is a good idea to have a backup solution in place to establish a secure connection between your networks. Cloud VPN is a secure and reliable solution for establishing a connection between your on-premises network and GCP. It uses a virtual private network (VPN) tunnel to securely connect the networks, and it is a good backup option if Dedicated Interconnect fails.

upvoted 8 times

omermahgoub 3 years, 8 months ago

The correct answer is B. Verify that Dedicated Interconnect can replicate files to GCP. Verify that Cloud VPN can establish a secure connection between your networks if Dedicated Interconnect fails. Dedicated Interconnect is a connection that provides a private, dedicated connection between your on-premises network and GCP over a Google-owned network. It is a secure and reliable option for connecting your on-premises network to GCP.

You can use it to replicate files to GCP as a part of your disaster recovery plan. If Dedicated Interconnect fails for any reason, it is a good idea to have a backup solution in place to establish a secure connection between your networks. Cloud VPN is a secure and reliable solution for establishing a connection between your on-premises network and GCP. It uses a virtual private network (VPN) tunnel to securely connect the networks, and it is a good backup option if Dedicated Interconnect fails.The Transfer Appliance is a physical storage device that you can use to transfer large amounts of data from your on-premises storage to GCP. It is not a connection option and cannot be used to establish a secure connection between your on-premises network and GCP. Therefore, the options C and D are not correct.Why not A? Is Cloud VPN better than Direct Peering in this scenario?

upvoted 8 times

omermahgoub 3 years, 8 months ago

The Transfer Appliance is a physical storage device that you can use to transfer large amounts of data from your on-premises storage to GCP. It is not a connection option and cannot be used to establish a secure connection between your on-premises network and GCP. Therefore, the options C and D are not correct.

upvoted 1 times

stefanop 2 years, 8 months ago

Why not A? Is Cloud VPN better than Direct Peering in this scenario?

upvoted 2 times

megumin 3 years, 9 months ago

B is ok

upvoted 1 times

zr79 3 years, 10 months ago

For DR with Google Cloud and on-prem use Dedicated Interconnect with HA VPN

upvoted 1 times

AzureDP900 3 years, 10 months ago

B is right without any second thought. Question is straight forward.

upvoted 1 times

Rajeev26 3 years, 10 months ago

Transfer appliance you need to carry to GCP center like water bottle :)

upvoted 1 times

abirroy 3 years, 11 months ago

Verify that Dedicated Interconnect can replicate files to GCP. Verify that direct peering can establish a secure connection between your networks if Dedicated Interconnect fails.

upvoted 1 times

alexandercamachop 3 years, 11 months ago

It is definitely B 1. Interconnect is the first option so that is right. 2. Eliminates A, since Direct Peering is not supported in GCP, the option is Google Cloud VPN connection to onpremises site.GCP supports direct peering in 100 locations

upvoted 3 times

alexandercamachop 3 years, 11 months ago

It is definitely B 1. Interconnect is the first option so that is right. 2. Eliminates A, since Direct Peering is not supported in GCP, the option is Google Cloud VPN connection to onpremises site.

upvoted 3 times

BeCalm 3 years, 5 months ago

GCP supports direct peering in 100 locations

upvoted 2 times

DrishaS4 4 years ago

Transfer appliance is a physical appliance for transferring huge bulk of data. does not fit into disaster recovery testing

upvoted 1 times

Matalf 4 years ago

only B have redundacy

upvoted 1 times

Question #81 *Topic 1*

Your company operates nationally and plans to use GCP for multiple batch workloads, including some that are not time-critical. You also need to use GCP services that are HIPAA-certified and manage service costs.
How should you design to meet Google best practices?

A. Provision preemptible VMs to reduce cost. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.

B. Provision preemptible VMs to reduce cost. Disable and then discontinue use of all GCP services and APIs that are not HIPAA-compliant.

C. Provision standard VMs in the same region to reduce cost. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.

D. Provision standard VMs to the same region to reduce cost. Disable and then discontinue use of all GCP services and APIs that are not HIPAA-compliant.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
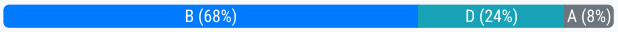

Eroc Highly Voted 6 years, 9 months ago
Disabling and then discontinuing allows you to see the effects of not using the APIs, so you can gauge (check) alternatives. So that leaves B and D as viable answers. The question says only some are not time-critical which implies others are... this means preemptible VMs are good because they will secure a spot for scaling when needed. So I'm also going to choose B.If others are time-critical, preemtible does not fit. Answer is D.No mention of others in the question. In an exam it's important to not being in individual assumptions and focus on the information in question. Key word here is "not time-critical"agree otherwise answer goes to non-preemtible VM'sVer well said..."In an exam it's important to not being in individual assumptions and focus on the information in question. Key word here is "not time-critical""And the others are not spoken about. By taking the question just by the context that it sets, preemptible is what I choose. So it's B according to me.Correctly explainedcorrectWhy you decided to emphasize on key word "not time-critical" but not" "operates nationally"?
upvoted 44 times

Eroc Highly Voted 6 years, 9 months ago
Disabling and then discontinuing allows you to see the effects of not using the APIs, so you can gauge (check) alternatives. So that leaves B and D as viable answers. The question says only some are not time-critical which implies others are... this means preemptible VMs are good because they will secure a spot for scaling when needed. So I'm also going to choose B.
upvoted 44 times

Musk 6 years ago
If others are time-critical, preemtible does not fit. Answer is D.
upvoted 11 times

Musk 6 years ago
If others are time-critical, preemtible does not fit. Answer is D.No mention of others in the question. In an exam it's important to not being in individual assumptions and focus on the information in question. Key word here is "not time-critical"agree otherwise answer goes to non-preemtible VM'sVer well said..."In an exam it's important to not being in individual assumptions and focus on the information in question. Key word here is "not time-critical""And the others are not spoken about. By taking the question just by the context that it sets, preemptible is what I choose. So it's B according to me.
upvoted 11 times

army234 5 years, 4 months ago
No mention of others in the question. In an exam it's important to not being in individual assumptions and focus on the information in question. Key word here is "not time-critical"
upvoted 7 times

army234 5 years, 4 months ago
No mention of others in the question. In an exam it's important to not being in individual assumptions and focus on the information in question. Key word here is "not time-critical"agree otherwise answer goes to non-preemtible VM'sVer well said..."In an exam it's important to not being in individual assumptions and focus on the information in question. Key word here is "not time-critical""
upvoted 7 times

AzureDP900 3 years, 10 months ago

agree otherwise answer goes to non-preemtible VM's

upvoted 2 times

Sur_Nikki 3 years, 3 months ago

Ver well said..."In an exam it's important to not being in individual assumptions and focus on the information in question. Key word here is "not time-critical""

upvoted 2 times

Darahaas 5 years, 11 months ago

And the others are not spoken about. By taking the question just by the context that it sets, preemptible is what I choose. So it's B according to me.

upvoted 3 times

Sur_Nikki 3 years, 3 months ago

Correctly explained

upvoted 1 times

AzureDP900 3 years, 10 months ago

correct

upvoted 1 times

netizens 2 years, 6 months ago

Why you decided to emphasize on key word "not time-critical" but not" "operates nationally"?

upvoted 1 times

Karna 6 years, 1 month ago

They say that some (not all) of the Batch workloads are not time critical which implies that there are time critical Batch workloads for which Preemptible VMs are not appropriate, so going with D as the answer

upvoted 20 times

CyberPeter Most Recent 6 months, 2 weeks ago

Same than B without disabling the services. With not use them is enought :). Disabling them needs extra effort

upvoted 1 times

JoseCloudEng1994 11 months, 3 weeks ago

Keywords "Batch" and "Not Time Sensitive". So its got to be preemptible VMs

upvoted 1 times

3a7557a 1 year, 6 months ago

. Provision standard VMs to the same region to reduce cost. Disable and then discontinue use of all GCP services and APIs that are not HIPAA-compliant. you can't use all premeptible, because only some are not time senstive, rest are time critical.. cash... D. keyword google best practice... if you have any time senstive u have to have some standard VMI.

upvoted 1 times

desertlotus1211 1 year, 8 months ago

The answer is A. Though some workloads are critical, the question is really asking about saving costs on the non-critical workloads. it's already understood that some workloads will incur a cost. So 'how' do we save in general?? Use preemptible VMs.

upvoted 1 times

nareshthumma 1 year, 9 months ago

Answer is B

upvoted 1 times

Ponchi14 2 years, 1 month ago

D is the correct answer and for the following regions: 1. Deploying a read replica and then manually failing over by stopping the current cloud SQL instance will only impact the authentication layer, which is what we're looking to test. 2. Stopping all VM's won't have an impact on the authentication layer testing becasuse Cloud SQL is a PaaS service, you cannot just turn off the VM's

upvoted 1 times

Sephethus 2 years, 2 months ago

What is the difference between A and B other than adding the obvious that you'd disable them if you discontinue using them. This is the most obnoxiously confusing question I've ever read and someone else pointed out that *some* are not time critical which implies others are time critical.

upvoted 1 times

JaimeMS 2 years, 2 months ago

If the question would have said: "The batch are not time-critical", then Option B with preentible VMs. BUT, it clearly points that only SOME are not time-critical". Option D is the only one that satisfies all conditions

upvoted 1 times

thewalker 2 years, 9 months ago

D As per the requirement, some of the batches are not time critical - which means some are critical. Choosing preemptible VMs may mean time critical batches may be affected. Cost effective solutions should not come at cost of requirements. Disabling and then discontinuing allows you to see the effects of not using the APIs, so you can gauge (check) alternatives.

upvoted 1 times

Arun_m_123 2 years, 10 months ago

When it is "some" batches are non-critical - when we want to focus on price reduction and if it is batch apps, then we can definitely choose "Pre-emptible VMs". Also non-compliant APIs needs to be disabled for sure (APIs can be enabled if there is a need). Putting altogether, B is the right answer

upvoted 2 times

Net50 3 years ago

Selected Answer: D

I think the answer should be D. some of the batches are not time critical which means some are. Choosing premtible VMs may mean time critical batches may be affected in some cases. Even though the solution needs to be cost effective, it should not come at cost of requirements. Hence D

upvoted 3 times

Net50 3 years ago

I think the answer should be D. some of the batches are not time critical which means some are. Choosing premtible VMs may mean time critical batches may be affected in some cases. Even though the solution needs to be cost effective, it should not come at cost of requirements. Hence DPremptible VMs is better for batch workloads which are critical and not critical since it takes bit longer to complete the load when some VMs are preempted . It greatly reduces cost with using preempted vms for batch workloads. B is the better answer

upvoted 3 times

Diwz 2 years, 4 months ago

Premptible VMs is better for batch workloads which are critical and not critical since it takes bit longer to complete the load when some VMs are preempted . It greatly reduces cost with using preempted vms for batch workloads. B is the better answer

upvoted 1 times

Sur_Nikki 3 years, 3 months ago

B looks good to me

upvoted 1 times

8d31d36 3 years, 6 months ago

To design a GCP solution that meets Google's best practices for operating nationally with multiple batch workloads, including some that are not time-critical, and using HIPAA-compliant services while managing service costs, you should provision standard VMs in the same region to reduce cost, and use GCP services that are HIPAA-compliant as needed. Therefore, the correct option is C. Preemptible VMs can provide cost savings, but they are not recommended for workloads that are not time-critical, as they may be interrupted at any time. Provisioning standard VMs in the same region will provide better performance and stability, and can still be cost-effective by using features such as sustained-use discounts and committed use discounts.

upvoted 1 times

megumin 3 years, 9 months ago

B is ok

upvoted 1 times

minmin2020 3 years, 10 months ago

Assumption here is that cost is more important that the time critical batches, therefore use preemptible instances. Disable and discontinue is a better option as it gives you opportunity to see the impact before blasting any APIs or services that are not certified.

upvoted 4 times

Question #82 *Topic 1*

Your customer wants to do resilience testing of their authentication layer. This consists of a regional managed instance group serving a public REST API that reads from and writes to a Cloud SQL instance.
What should you do?

A. Engage with a security company to run web scrapers that look your for users' authentication data om malicious websites and notify you if any is found.

B. Deploy intrusion detection software to your virtual machines to detect and log unauthorized access.

C. Schedule a disaster simulation exercise during which you can shut off all VMs in a zone to see how your application behaves.

D. Configure a read replica for your Cloud SQL instance in a different zone than the master, and then manually trigger a failover while monitoring KPIs for our REST API.

Suggested Answer: *C*
Official Answer: C

Community vote distribution

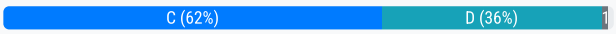

Kri_2525 Highly Voted 6 years, 8 months ago
As per google documentation(https://cloud.google.com/solutions/scalable-and-resilient-apps) answer is C. C: A well-designed application should scale seamlessly as demand increases and decreases, and be resilient enough to withstand the loss of one or more compute resources. Resilience: designed to withstand the unexpected A highly-available, or resilient, application is one that continues to function despite expected or unexpected failures of components in the system. If a single instance fails or an entire zone experiences a problem, a resilient application remains fault tolerant—continuing to function and repairing itself automatically if necessary. Because stateful information isn't stored on any single instance, the loss of an instance—or even an entire zone—should not impact the application's performance.Shutting off all VMs in a zone is not good approach for testing of authenticationIt is good approach. Netflix does it. Dont forget that MIG is regional, so probably there will be some machines in the rest AZs.You're not testing *authentication*, you're testing *the resilience of the authentication layer*. "A resilient app is one that continues to function despite failures of system components" (https://cloud.google.com/architecture/scalable-and-resilient-apps#resilience_designing_to_withstand_failures) - such as shutting down all VMs in a zone.Agree, Chaos testing is to shutdown random instances.yes chaos testing is industry standardNote: resilience and high availability are not the same, though they are closely related. High availability is a part of an overall resilience strategy. High availability focuses on keeping a system up and running with minimal downtime. It is primarily concerned with eliminating single points of failure by using redundancy. The goal is to ensure a system is accessible and operational as close to 100% of the time as possible. Key Question: Is the system available? How it's achieved: Through redundant components, load balancing, and a
upvoted 66 times

Kri_2525 Highly Voted 6 years, 8 months ago
As per google documentation(https://cloud.google.com/solutions/scalable-and-resilient-apps) answer is C.

C: A well-designed application should scale seamlessly as demand increases and decreases, and be resilient enough to withstand the loss of one or more compute resources. Resilience: designed to withstand the unexpected A highly-available, or resilient, application is one that continues to function despite expected or unexpected failures of components in the system. If a single instance fails or an entire zone experiences a problem, a resilient application remains fault tolerant—continuing to function and repairing itself automatically if necessary. Because stateful information isn't stored on any single instance, the loss of an instance—or even an entire zone—should not impact the application's performance.
upvoted 66 times

Jack_in_Large 6 years, 2 months ago
Shutting off all VMs in a zone is not good approach for testing of authentication
upvoted 6 times

Jack_in_Large 6 years, 2 months ago
Shutting off all VMs in a zone is not good approach for testing of authenticationIt is good approach. Netflix does it. Dont forget that MIG is regional, so probably there will be some machines in the rest AZs.You're not testing *authentication*, you're testing *the resilience of the authentication layer*. "A resilient app is one that continues to function despite failures of system components" (https://cloud.google.com/architecture/scalable-and-resilient-apps#resilience_designing_to_withstand_failures) - such as shutting down all VMs in a zone.Agree, Chaos testing is to shutdown random instances.yes chaos testing is industry standard

upvoted 6 times

hitmax87 Most Recent 2 months, 2 weeks ago

It is good approach. Netflix does it. Dont forget that MIG is regional, so probably there will be some machines in the rest AZs.

upvoted 1 times

vartiklis 4 years, 8 months ago

You're not testing *authentication*, you're testing *the resilience of the authentication layer*. "A resilient app is one that continues to function despite failures of system components" (https://cloud.google.com/architecture/scalable-and-resilient-apps#resilience_designing_to_withstand_failures) - such as shutting down all VMs in a zone.

upvoted 20 times

elaineshi 4 years, 2 months ago

Agree, Chaos testing is to shutdown random instances.

upvoted 4 times

elaineshi 4 years, 2 months ago

Agree, Chaos testing is to shutdown random instances.yes chaos testing is industry standard

upvoted 4 times

heretolearnazure 2 years, 12 months ago

yes chaos testing is industry standard

upvoted 3 times

JoseCloudEng1994 11 months, 3 weeks ago

Note: resilience and high availability are not the same, though they are closely related. High availability is a part of an overall resilience strategy. High availability focuses on keeping a system up and running with minimal downtime. It is primarily concerned with eliminating single points of failure by using redundancy. The goal is to ensure a system is accessible and operational as close to 100% of the time as possible. Key Question: Is the system available? How it's achieved: Through redundant components, load balancing, and automatic failover. If one component fails, a backup is immediately ready to take over.Resilience is a broader concept. It is the ability of a system to not only survive a failure but also to adapt to and recover quickly from adverse conditions or disruptions. It's about how the system handles unforeseen failures and operates even in a degraded state. Key Question: Can the system recover from failure? How it's achieved: Through strategies like self-healing, graceful degradation, and a robust disaster recovery plan. A resilient system may experience a temporary dip in performance, but it can quickly bounce back to an effective operational state.

upvoted 1 times

JoseCloudEng1994 11 months, 3 weeks ago

Note: resilience and high availability are not the same, though they are closely related. High availability is a part of an overall resilience strategy. High availability focuses on keeping a system up and running with minimal downtime. It is primarily concerned with eliminating single points of failure by using redundancy. The goal is to ensure a system is accessible and operational as close to 100% of the time as possible.

Key Question: Is the system available?

How it's achieved: Through redundant components, load balancing, and automatic failover. If one component fails, a backup is immediately ready to take over.

upvoted 1 times

JoseCloudEng1994 11 months, 3 weeks ago

Resilience is a broader concept. It is the ability of a system to not only survive a failure but also to adapt to and recover quickly from adverse conditions or disruptions. It's about how the system handles unforeseen failures and operates even in a degraded state. Key Question: Can the system recover from failure? How it's achieved: Through strategies like self-healing, graceful degradation, and a robust disaster recovery plan. A resilient system may experience a temporary dip in performance, but it can quickly bounce back to an effective operational state.

upvoted 1 times

KouShikyou 6 years, 9 months ago

Since the question is asking to do a resilience testing, I prefer C.Resilience testing of the "Authentication Layer", not the "Application". So the answer is B.

upvoted 16 times

KouShikyou 6 years, 9 months ago

Since the question is asking to do a resilience testing, I prefer C.

upvoted 16 times

Darahaas 5 years, 11 months ago

Resilience testing of the "Authentication Layer", not the "Application". So the answer is B.

upvoted 4 times

Pime13 6 months ago

Selected Answer: C

Regional MIG Resilience: A regional managed instance group is designed to survive zonal outages by automatically redistributing or recreating instances in healthy zones. Shutting off all VMs in a zone is a standard "chaos engineering" practice to validate that your application remains available and that your autoscaling/self-healing configurations work as expected under stress. Google-Recommended Practice: Google explicitly provides scripts and metadata methods to simulate zonal failures for regional MIGs to test if they are correctly overprovisioned and capable of surviving such events. .

upvoted 1 times

Pime13 6 months ago

Regional MIG Resilience: A regional managed instance group is designed to survive zonal outages by automatically redistributing or recreating instances in healthy zones. Shutting off all VMs in a zone is a standard "chaos engineering" practice to validate that your application remains available and that your autoscaling/self-healing configurations work as expected under stress. Google-Recommended Practice: Google explicitly provides scripts and metadata methods to simulate zonal failures for regional MIGs to test if they are correctly overprovisioned and capable of surviving such events. .Targeting the Architecture: Since your application uses a regional MIG, testing its ability to handle a full zone being offline is the most comprehensive way to validate the "resilience" of the authentication layer. not D: While this tests the database's High Availability (HA) failover, a regional managed instance group is primarily an infrastructure-level tool. Testing the VMs themselves (Option C) provides a broader validation of the entire application stack's ability to reroute traffic and maintain service during a disaster. In a true resilience test, you want to see how the application handles the combined loss of compute and potential database reconnects.

upvoted 1 times

Pime13 6 months ago

Targeting the Architecture: Since your application uses a regional MIG, testing its ability to handle a full zone being offline is the most comprehensive way to validate the "resilience" of the authentication layer. not D: While this tests the database's High Availability (HA) failover, a regional managed instance group is primarily an infrastructure-level tool. Testing the VMs themselves (Option C) provides a broader validation of the entire application stack's ability to reroute traffic and maintain service during a disaster. In a true resilience test, you want to see how the application handles the combined loss of compute and potential database reconnects.

upvoted 1 times

DM____1111 1 year, 1 month ago

This approach directly tests how your application's authentication layer behaves when its database experiences a failover event, which is crucial for resilience. C is not correct because while shutting off all VMs in a zone is a valid and comprehensive disaster recovery test for a full zonal outage, it's a very broad test. The question specifically asks about testing the authentication layer's resilience and its reliance on Cloud SQL. Directly testing the Cloud SQL failover is more focused and efficient for this particular goal.

upvoted 1 times

f3e7abb 1 year, 2 months ago

Resilience in this case means at least be able to support a failure of a zone, so considering that the MIG is regional, so it can withstand a zonal failure, the only other thing is to have a failover strategy for Cloud SQL and triggering it.

upvoted 1 times

MrAZ105 1 year, 3 months ago

D is correct

upvoted 1 times

khadar2001 1 year, 4 months ago

The requirement is on authentication layer so selected answer B. This is confirmed from GCP Practice exam.

upvoted 1 times

david_tay 1 year, 5 months ago

answer is D, check in Gemini to understand why.

upvoted 1 times

Ekramy_Elnaggar 1 year, 9 months ago

D. is not correct as this tests the resilience of the database (Cloud SQL) but not necessarily the authentication layer. The authentication layer might have separate components or dependencies that need to be tested under failure conditions.

upvoted 2 times

nareshthumma 1 year, 9 months ago

Agree with C

upvoted 2 times

wooyourdaddy 1 year, 10 months ago

Option C, which involves scheduling a disaster simulation exercise to shut off all VMs in a zone, is indeed a strong choice for resilience testing. This approach helps you understand how your application behaves under failure conditions and ensures that your system can handle unexpected disruptions. However, Option D is also highly relevant. Configuring a read replica for your Cloud SQL instance in a different zone and manually

triggering a failover while monitoring KPIs for your REST API directly tests the resilience of your database layer. This can provide valuable insights into the failover process and the impact on your application's performance and availability. Both options have their merits, but if the primary goal is to test the resilience of the authentication layer specifically, Option D might be more targeted and effective.

upvoted 2 times

hitmax87 2 years, 3 months ago

C is correct. It is not D because you are not designing system, your goal is testing existed system

upvoted 4 times

666Amitava666 2 years, 3 months ago

Chaos testing

upvoted 4 times

activist 2 years, 4 months ago

I choose Answer C https://cloud.google.com/sql/docs/mysql/replication This URL states "Read replicas are read-only; you cannot write to them. The read replica processes queries, read requests, and analytics traffic, thus reducing the load on the primary instance." "Note: Read replicas do not provide failover capability. To provide failover capability for an instance, see Configuring an instance for high availability." "As a best practice, put read replicas in a different zone than the primary instance when you use HA on your primary instance. This practice ensures that read replicas continue to operate when the zone that contains the primary instance has an outage. See the Overview of high availability for more information."

upvoted 3 times

santoshchauhan 2 years, 4 months ago

Testing Database Resilience: By setting up a read replica in a different zone and triggering a manual failover, you simulate a failure of the primary database. This allows you to assess how well your authentication layer and the overall application cope with the loss of the primary database. Monitoring Performance and Availability: During the failover, monitoring key performance indicators (KPIs) for your REST API will give you insights into how the application's performance and availability are impacted. This helps in identifying potential bottlenecks and areas for improvement in your resilience strategy. Ensuring Data Continuity: A read replica ensures data continuity and minimizes downtime, which is critical for an authentication system. The replica will take over as the primary database during the failover, ensuring that the authentication service remains functional.

upvoted 1 times

Rehamss 2 years, 5 months ago

D is okay

upvoted 1 times

Teckexam 2 years, 6 months ago

Authentication layer resiliency can be covered as part of overall application resiliency testing. Option C is asking to use read replica which is not useful in case of testing resiliency in case of failure

upvoted 3 times

Question #83 *Topic 1*

Your BigQuery project has several users. For audit purposes, you need to see how many queries each user ran in the last month. What should you do?

A. Connect Google Data Studio to BigQuery. Create a dimension for the users and a metric for the amount of queries per user.

B. In the BigQuery interface, execute a query on the JOBS table to get the required information.

C. Use 'bq show' to list all jobs. Per job, use 'bq ls' to list job information and get the required information.

D. Use Cloud Audit Logging to view Cloud Audit Logs, and create a filter on the query operation to get the required information.

Suggested Answer: *C*
Official Answer: C

Community vote distribution

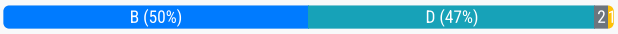

Googler2 Highly Voted 6 years, 4 months ago
D- reasons: 1.-Cloud Audit Logs maintains audit logs for admin activity, data access and system events. BIGQUERY is automatically send to cloud audit log functionality. 2.- In the filter you can filter relevant BigQuery Audit messages, you can express filters as part of the export https://cloud.google.com/logging/docs/audit https://cloud.google.com/bigquery/docs/reference/auditlogs#ids https://cloud.google.com/bigquery/docs/reference/auditlogs#auditdata_examplesAnswer is DD is the right as you can get the monthly view of the query usage across all the users and projects for auditing purpose. C does need appropriate permission to see the detail level data. Monthly view is tough to get directly from the bq ls or bq show commands.
upvoted 54 times

Googler2 Highly Voted 6 years, 4 months ago
D- reasons: 1.-Cloud Audit Logs maintains audit logs for admin activity, data access and system events. BIGQUERY is automatically send to cloud audit log functionality. 2.- In the filter you can filter relevant BigQuery Audit messages, you can express filters as part of the export

https://cloud.google.com/logging/docs/audit https://cloud.google.com/bigquery/docs/reference/auditlogs#ids https://cloud.google.com/bigquery/docs/reference/auditlogs#auditdata_examples
upvoted 54 times

heretolearnazure 2 years, 12 months ago
Answer is D
upvoted 3 times

GooglecloudArchitect 6 years ago
D is the right as you can get the monthly view of the query usage across all the users and projects for auditing purpose. C does need appropriate permission to see the detail level data. Monthly view is tough to get directly from the bq ls or bq show commands.
upvoted 10 times

Zarmi 6 years, 3 months ago
Answer is D: https://cloud.google.com/bigquery/docs/reference/auditlogs#example_query_cost_breakdown_by_identity
upvoted 27 times

Zarmi 6 years, 3 months ago
Answer is D: https://cloud.google.com/bigquery/docs/reference/auditlogs#example_query_cost_breakdown_by_identityNailed itNo mention about exporting to bq
upvoted 27 times

BobbyFlash 4 years, 7 months ago
Nailed it
upvoted 3 times

ErenYeager 3 years, 9 months ago
No mention about exporting to bq
upvoted 2 times

0d8df8b 2 months, 2 weeks ago

The task is to find out how many queries each user executed over the past month. BigQuery stores job metadata in the INFORMATION_SCHEMA.JOBS views.

upvoted 2 times

Johnb9682 4 months, 1 week ago

Selected Answer: D

Answer is D. Reason: Cloud Audit Logs are the official and definitive source for audit purposes.

upvoted 1 times

Johnb9682 4 months, 1 week ago

Answer is D. Reason: Cloud Audit Logs are the official and definitive source for audit purposes.Audit Logs is used to identify who and when performed actions on GCP infrastructure. Correct answer is B

upvoted 1 times

hitmax87 Most Recent 2 months, 2 weeks ago

Audit Logs is used to identify who and when performed actions on GCP infrastructure. Correct answer is B

upvoted 1 times

acloudboy 5 months, 1 week ago

Answer is D. Reason: Cloud Audit Logs are the official and definitive source for audit purposes.

upvoted 1 times

Pime13 6 months ago

Selected Answer: D

Comprehensive Tracking: BigQuery automatically integrates with Cloud Audit Logs, which record every API call and operation, including query execution (job creation). This provides a reliable, tamper-proof history of all user activity required for formal audits.

Granular Filtering: Within the Cloud Logging interface, you can use the Log Explorer to create specific filters. By filtering for the google.cloud.bigquery.v2.JobService.Insert method and the query operation over a specific time range (the last month), you can accurately group and count the total number of queries initiated by each individual user identity (principalEmail).

upvoted 1 times

Pime13 6 months ago

Comprehensive Tracking: BigQuery automatically integrates with Cloud Audit Logs, which record every API call and operation, including query execution (job creation). This provides a reliable, tamper-proof history of all user activity required for formal audits. Granular Filtering: Within the Cloud Logging interface, you can use the Log Explorer to create specific filters. By filtering for the google.cloud.bigquery.v2.JobService.Insert method and the query operation over a specific time range (the last month), you can accurately group and count the total number of queries initiated by each individual user identity (principalEmail).Scale and Retention: Unlike manual listing commands or direct table queries, Cloud Audit Logs are designed to handle millions of events and offer customizable retention periods suitable for long-term compliance and historical analysis.

upvoted 1 times

Pime13 6 months ago

Scale and Retention: Unlike manual listing commands or direct table queries, Cloud Audit Logs are designed to handle millions of events and offer customizable retention periods suitable for long-term compliance and historical analysis.

upvoted 1 times

CyberPeter 6 months, 2 weeks ago

BigQuery keeps metadata about all jobs (including queries) in its job history tables. You can query these (for example via INFORMATION_SCHEMA.JOBS*) to: Filter to the last month Group by user (principal_email) Count the number of QUERY jobs This is the standard, auditable, and exam-expected way to answer questions like “how many queries did each user run?”

upvoted 1 times

KCjoe 8 months ago

Google Cloud's BigQuery service automatically logs metadata about every job run in your project into system tables for up to 60 days. This makes auditing query usage straightforward using standard SQL queries.

upvoted 3 times

ast3citos 10 months ago

Apparently while D is historizally correct and gets you the data you are looking for, B is the more modern approach introduced by Google to query for jobs metadata.

upvoted 3 times

elon_musk 11 months, 2 weeks ago

simpest and cheapest

upvoted 2 times

JoseCloudEng1994 11 months, 3 weeks ago

Selected Answer: B

For audit and security purposes, the definitive source of truth for all activity in a Google Cloud project is Cloud Audit Logs. BigQuery automatically generates audit logs that record a wide range of activities, including when a user runs a query.

This approach is the most reliable because:

It's the official, comprehensive audit trail for your project.

You can filter logs by methodName: "google.cloud.bigquery.v2.JobService.Query" to specifically target query jobs.

The logs contain metadata about the user who initiated the job, the timestamp, and other details necessary for your audit.

upvoted 2 times

JoseCloudEng1994 11 months, 3 weeks ago

For audit and security purposes, the definitive source of truth for all activity in a Google Cloud project is Cloud Audit Logs. BigQuery automatically generates audit logs that record a wide range of activities, including when a user runs a query. This approach is the most reliable because: It's the official, comprehensive audit trail for your project. You can filter logs by methodName: "google.cloud.bigquery.v2.JobService.Query" to specifically target query jobs. The logs contain metadata about the user who initiated the job, the timestamp, and other details necessary for your audit.B. In the BigQuery interface, execute a query on the JOBS table to get the required information. □ This option is close but refers to a non-existent "JOBS table" in the traditional sense. The correct way to query job metadata within BigQuery itself is to use the INFORMATION_SCHEMA.JOBS view. This view contains metadata about all jobs (queries, loads, etc.) run in the project. You can query this view to get the number of queries per user. However, this is for performance analysis and monitoring, and while it can be used for audits, Cloud Audit Logs are the official and definitive source for audit purposes.

upvoted 2 times

JoseCloudEng1994 11 months, 3 weeks ago

B. In the BigQuery interface, execute a query on the JOBS table to get the required information. □ This option is close but refers to a non-existent "JOBS table" in the traditional sense. The correct way to query job metadata within BigQuery itself is to use the INFORMATION_SCHEMA.JOBS view. This view contains metadata about all jobs (queries, loads, etc.) run in the project. You can query this view to get the number of queries per user. However, this is for performance analysis and monitoring, and while it can be used for audits, Cloud Audit Logs are the official and definitive source for audit purposes.

upvoted 1 times

rajrocks171 1 year, 2 months ago

Option D. Cloud audit logs track all BiqQuery jobs including who ran each query, when and how. You can filter logs by operation type and group by user identity. Not Option B - Unless you are a project admin, by default JOBS table shows jobs for the current user. Also, retention period here is 7 days.

upvoted 4 times

lokiinaction 1 year, 2 months ago

Audit log (exported to BQ) is source for audit trails and security, includes full SQL text, can be retained for very long periods if exported.

upvoted 3 times

markware 1 year, 3 months ago

This is mentioned in Engineer course from Google

upvoted 3 times

francisco94 1 year, 4 months ago

I agree B could be the solution, but the best option is D. This is the correct and scalable approach: Cloud Audit Logs capture who ran what, including queries. You can filter on methodName = "jobservice.jobcompleted" and analyze logs in Logs Explorer or export to BigQuery for querying. Supports organization-wide, project-wide, and per-user visibility.

upvoted 3 times

Mikeliz 1 year, 5 months ago

B is the better answer. You get more logs relating to BigQuery jobs from BigQuery than you get from Cloud Audit Logs

upvoted 2 times

david_tay 1 year, 5 months ago

B is correct. Go search in Gemini "what are the steps to execute a query on the JOBS table in Bigquery to see how many queries each user ran in the last month." to see how easy are the steps.

upvoted 2 times

Question #84 *Topic 1*

You want to automate the creation of a managed instance group. The VMs have many OS package dependencies. You want to minimize the startup time for new
VMs in the instance group.
What should you do?

A. Use Terraform to create the managed instance group and a startup script to install the OS package dependencies.

B. Create a custom VM image with all OS package dependencies. Use Deployment Manager to create the managed instance group with the VM image.

C. Use Puppet to create the managed instance group and install the OS package dependencies.

D. Use Deployment Manager to create the managed instance group and Ansible to install the OS package dependencies.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (100%)

crypt0 Highly Voted 5 years, 4 months ago
Why is it not answer B?It is.B is okIt is BB is the answer,
upvoted 42 times

crypt0 Highly Voted 5 years, 4 months ago
Why is it not answer B?
upvoted 42 times

Jos 5 years, 1 month ago
It is.B is ok
upvoted 10 times

Jos 5 years, 1 month ago
It is.
upvoted 10 times

tartar 4 years, 6 months ago
B is ok
upvoted 12 times

nitinz 3 years, 11 months ago
It is B
upvoted 4 times

kumarp6 4 years, 3 months ago
B is the answer,
upvoted 6 times

JoeShmoe 5 years, 3 months ago
B- minimal start time means a pre-baked golden image
upvoted 21 times

omermahgoub 2 years, 1 month ago
The correct answer is B. Create a custom VM image with all OS package dependencies. Use Deployment Manager to create the managed instance group with the VM image. Managed instance groups are a way to manage a group of Compute Engine instances as a single entity. If you want to automate the creation of a managed instance group, you can use tools such as Terraform, Deployment Manager, or Puppet to automate the process. To minimize the startup time for new VMs in the instance group, you should create a custom VM image with all of the OS package dependencies pre-installed. This will allow you to create new VMs from the custom image, which will significantly reduce the startup time compared to installing the dependencies on each VM individually. You can then use Deployment Manager to create the managed instance group with the custom VM image.Option A, using Terraform to create the managed instance group and a startup script to install the OS package dependencies, would not minimize the startup time for new VMs in the instance group. Option C, using Puppet to create the managed instance group and install the OS package dependencies, would not minimize the startup time for new VMs in the instance group. Option D, using Deployment Manager to create the

managed instance group and Ansible to install the OS package dependencies, would not minimize the startup time for new VMs in the instance group.

upvoted 15 times

omermahgoub 2 years, 1 month ago

The correct answer is B. Create a custom VM image with all OS package dependencies. Use Deployment Manager to create the managed instance group with the VM image.

Managed instance groups are a way to manage a group of Compute Engine instances as a single entity. If you want to automate the creation of a managed instance group, you can use tools such as Terraform, Deployment Manager, or Puppet to automate the process.

To minimize the startup time for new VMs in the instance group, you should create a custom VM image with all of the OS package dependencies pre-installed. This will allow you to create new VMs from the custom image, which will significantly reduce the startup time compared to installing the dependencies on each VM individually. You can then use Deployment Manager to create the managed instance group with the custom VM image.

upvoted 15 times

omermahgoub Most Recent 2 years, 1 month ago

Option A, using Terraform to create the managed instance group and a startup script to install the OS package dependencies, would not minimize the startup time for new VMs in the instance group. Option C, using Puppet to create the managed instance group and install the OS package dependencies, would not minimize the startup time for new VMs in the instance group. Option D, using Deployment Manager to create the managed instance group and Ansible to install the OS package dependencies, would not minimize the startup time for new VMs in the instance group.

upvoted 5 times

megumin 2 years, 3 months ago

B is ok

upvoted 1 times

AzureDP900 2 years, 4 months ago

B. Create a custom VM image with all OS package dependencies. Use Deployment Manager to create the managed instance group with the VM image.

upvoted 1 times

minmin2020 2 years, 4 months ago

B will reduce the startup time

upvoted 1 times

DrishaS4 2 years, 6 months ago

B- minimal start time means a pre-baked golden image

upvoted 4 times

BigSteveO 2 years, 7 months ago

As someone who works on Terraform. It may not be Googles best practice, even though it's built in just need to be initialized. But it is the easiest way to build and restructure infrastructure with a simple line of code change and a quick shell command to apply terraform. I mean B would work. But it doesn't include the start-up script for the OS dependencies to be loaded. ?>?>? Any feedback?

upvoted 1 times

BigSteveO 2 years, 7 months ago

As someone who works on Terraform. It may not be Googles best practice, even though it's built in just need to be initialized. But it is the easiest way to build and restructure infrastructure with a simple line of code change and a quick shell command to apply terraform. I mean B would work. But it doesn't include the start-up script for the OS dependencies to be loaded. ?>?>? Any feedback?Start up scripts aren't need here as you're making a custom OS image with all OS package dependencies. Question is not asking for the easiest way, it's asking how to minimize VM startup times. Not having to run the startup scripts because it's baked into the image is how I understand and interpret this, therefore B.

upvoted 1 times

Ric350 2 years, 7 months ago

Start up scripts aren't need here as you're making a custom OS image with all OS package dependencies. Question is not asking for the easiest way, it's asking how to minimize VM startup times. Not having to run the startup scripts because it's baked into the image is how I understand and interpret this, therefore B.

upvoted 2 times

mv2000 2 years, 7 months ago

06/30/2022 Exam

upvoted 5 times

rogerlovato 3 years, 1 month ago

B is correct

upvoted 1 times

haroldbenites 3 years, 2 months ago

Go for B

upvoted 1 times

Godlike 3 years, 2 months ago

yes B is right

upvoted 2 times

vincy2202 3 years, 2 months ago

B is the right answer

upvoted 2 times

exam_war 3 years, 3 months ago

go with B. D: it involves so many other third software to configure/manage which makes build more complicated.

upvoted 1 times

MaxNRG 3 years, 3 months ago

B – create a custom VM instance image with all OS dependencies. Use Deployment Manager to create a MIG with the VM image. Read more about Public and Custom VM Images: https://cloud.google.com/compute/docs/images Custom images are available in your project only, they don't add cost to your VM instances, incur image storage cost (0.085$ GB/month) D – could be also an alternative (if to consider requirement to install dependencies in start up script). But, last sentence stresses on "minimize VM's start up time". So, B is fastest solution. Also, what is a point to use Ansible if you can complete same task via startup script of Deployment Manager. Ansible won't make this faster, but just will add 3rd party dependency.

upvoted 3 times

victory108 3 years, 9 months ago

B. Create a custom VM image with all OS package dependencies. Use Deployment Manager to create the managed instance group with the VM image.

upvoted 2 times

un 3 years, 9 months ago

B is correct

upvoted 1 times

Question #85 *Topic 1*

Your company captures all web traffic data in Google Analytics 360 and stores it in BigQuery. Each country has its own dataset. Each dataset has multiple tables.
You want analysts from each country to be able to see and query only the data for their respective countries.
How should you configure the access rights?

A. Create a group per country. Add analysts to their respective country-groups. Create a single group 'all_analysts', and add all country-groups as members. Grant the 'all_analysts' group the IAM role of BigQuery jobUser. Share the appropriate dataset with view access with each respective analyst country-group.

B. Create a group per country. Add analysts to their respective country-groups. Create a single group 'all_analysts', and add all country-groups as members. Grant the 'all_analysts' group the IAM role of BigQuery jobUser. Share the appropriate tables with view access with each respective analyst country-group.

C. Create a group per country. Add analysts to their respective country-groups. Create a single group 'all_analysts', and add all country-groups as members. Grant the 'all_analysts' group the IAM role of BigQuery dataViewer. Share the appropriate dataset with view access with each respective analyst country- group.

D. Create a group per country. Add analysts to their respective country-groups. Create a single group 'all_analysts', and add all country-groups as members. Grant the 'all_analysts' group the IAM role of BigQuery dataViewer. Share the appropriate table with view access with each respective analyst country-group.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (65%) C (35%)

Sebatian Highly Voted 6 years, 8 months ago
It should be A. The question requires that user from each country can only view a specific data set, so BQ dataViewer cannot be assigned at project level. Only A could limit the user to query and view the data that they are supposed to be allowed to.
upvoted 63 times

Sebatian Highly Voted 6 years, 8 months ago
It should be A. The question requires that user from each country can only view a specific data set, so BQ dataViewer cannot be assigned at project level. Only A could limit the user to query and view the data that they are supposed to be allowed to.Should be C. https://cloud.google.com/bigquery/docs/access-control#bigquery.dataViewer Data viewer role can be applied to a Table and a View. JobUser can be applied only at a Project level not at a Dataset level https://cloud.google.com/bigquery/docs/access-control#bigquery.jobUserincorrect, should be A, BigQuery Job User (roles/bigquery.jobUser) Provides permissions to run jobs, including queries, within the project.A is wrong
upvoted 63 times

jits1984 3 years, 4 months ago
Should be C.

https://cloud.google.com/bigquery/docs/access-control#bigquery.dataViewer

Data viewer role can be applied to a Table and a View.

JobUser can be applied only at a Project level not at a Dataset level

https://cloud.google.com/bigquery/docs/access-control#bigquery.jobUser
upvoted 12 times

jits1984 3 years, 4 months ago
Should be C. https://cloud.google.com/bigquery/docs/access-control#bigquery.dataViewer Data viewer role can be applied to a Table and a View. JobUser can be applied only at a Project level not at a Dataset level https://cloud.google.com/bigquery/docs/access-control#bigquery.jobUserincorrect, should be A, BigQuery Job User (roles/bigquery.jobUser) Provides permissions to run jobs, including queries, within the project.A is wrong

upvoted 12 times

jits1984 2 years, 11 months ago

incorrect, should be A, BigQuery Job User (roles/bigquery.jobUser)

Provides permissions to run jobs, including queries, within the project.

upvoted 3 times

jits1984 2 years, 11 months ago

incorrect, should be A, BigQuery Job User (roles/bigquery.jobUser) Provides permissions to run jobs, including queries, within the project.A is wrong

upvoted 3 times

RKS_2021 2 years, 10 months ago

A is wrong

upvoted 1 times

wk 6 years, 10 months ago

Should be C https://cloud.google.com/bigquery/docs/access-control#bigquery.dataViewer When applied to a dataset, dataViewer provides permissions to: Read the dataset's metadata and to list tables in the dataset. Read data and metadata from the dataset's tables. When applied at the project or organization level, this role can also enumerate all datasets in the project. Additional roles, however, are necessary to allow the running of jobs.Option C grant read permission to all datasets globally, which violated the request "You want analysts from each country to be able to see and query only the data for their respective countries" So the correct answer is A.https://cloud.google.com/bigquery/docs/access-control#bigquery.dataViewer "When applied to a dataset.." you can apply dataViewer role to a specific dataset.

upvoted 32 times

wk 6 years, 10 months ago

Should be C https://cloud.google.com/bigquery/docs/access-control#bigquery.dataViewer

When applied to a dataset, dataViewer provides permissions to:

Read the dataset's metadata and to list tables in the dataset. Read data and metadata from the dataset's tables. When applied at the project or organization level, this role can also enumerate all datasets in the project. Additional roles, however, are necessary to allow the running of jobs.

upvoted 32 times

Jack_in_Large 6 years, 3 months ago

Option C grant read permission to all datasets globally, which violated the request "You want analysts from each country to be able to see and query only the data for their respective countries" So the correct answer is A.https://cloud.google.com/bigquery/docs/access-control#bigquery.dataViewer "When applied to a dataset.." you can apply dataViewer role to a specific dataset.

upvoted 32 times

Jack_in_Large 6 years, 3 months ago

Option C grant read permission to all datasets globally, which violated the request "You want analysts from each country to be able to see and query only the data for their respective countries"

So the correct answer is A.

upvoted 32 times

BrunoTostes 4 years, 10 months ago

https://cloud.google.com/bigquery/docs/access-control#bigquery.dataViewer "When applied to a dataset.." you can apply dataViewer role to a specific dataset.

upvoted 9 times

cloud_rider Most Recent 1 year, 5 months ago

Answer should be C. As the lowest possible level to grant jobUser role is project, all the analyst will get access to all the datasets in the project. Whereas the lowest possible level to dataViewer is Table/View, hence the required restriction can be setup by this approach only.

upvoted 1 times

d83229d 1 year, 6 months ago

All provided options seem to be bad practice. If you grant a group of job user on the project level (The lowest level allowed), you essentially grant them the ability to query all counties datasets. However, you can allow data viewer at the resource level which is what makes the most sense. But specifically for this question, Job user on the project level is less bad than Data viewer on the project level since both let you run queries but data viewer also allows you to view the dataset. Hence , A

upvoted 2 times

venkee 1 year, 6 months ago

in answr 'A' theis sentence "Share the appropriate dataset with view access with each respective analyst country-group" is very cryptic. I assume this means assign "data viewer" role for the respective analyst groups for the country specific datasets

upvoted 1 times

Cloud_Architect_05 1 year, 8 months ago

Should be A. DataViewer: "When applied at the project or organization level, this role can also enumerate all datasets in the project. Additional roles, however, are necessary to allow the running of jobs." https://cloud.google.com/bigquery/docs/access-control

upvoted 1 times

nareshthumma 1 year, 9 months ago

C. Explanation: Each country should have its own group to manage access efficiently. This allows you to easily add or remove analysts from their respective groups. By adding analysts to their specific country groups, you can manage permissions in a way that aligns with their data access needs. This group will include all country groups. It simplifies the management of roles for all analysts collectively. The dataViewer role provides permission to view datasets and tables. This role allows analysts to read data without the ability to modify it, which is appropriate for your use case. Granting view access to the respective datasets for each country group ensures that analysts can only access data relevant to their country. This is crucial for maintaining data privacy and compliance. Why Other Options Are Less Suitable: Using BigQuery jobUser Role: The BigQuery jobUser role allows users to run jobs (like queries) but does not inherently grant access to view datasets or tables. This option would not effectively limit visibility to data by country.

upvoted 2 times

Diwz 2 years, 5 months ago

Selected Answer: C

It is C. Question says analyst should be able to see and query only the data for their respective countries. BigQueryDta viewer permission will allow only to read and query the table/view data

upvoted 3 times

Diwz 2 years, 5 months ago

It is C. Question says analyst should be able to see and query only the data for their respective countries. BigQueryDta viewer permission will allow only to read and query the table/view dataYou cant query with dataviewer. user with the roles/bigquery.dataViewer role has read-only access to datasets and tables but does not inherently have the permissions to run queries (which are considered jobs in BigQuery). The dataViewer role allows users to view dataset metadata and table contents but does not include the ability to create or execute jobs.The dataViewer role alone does not allow users to run queries. Analysts need the ability to run queries, which requires the jobUser role.

upvoted 3 times

awsgcparch 2 years ago

You cant query with dataviewer. user with the roles/bigquery.dataViewer role has read-only access to datasets and tables but does not inherently have the permissions to run queries (which are considered jobs in BigQuery). The dataViewer role allows users to view dataset metadata and table contents but does not include the ability to create or execute jobs.The dataViewer role alone does not allow users to run queries. Analysts need the ability to run queries, which requires the jobUser role.

upvoted 2 times

OrangeTiger 2 years, 6 months ago

Go with a.

upvoted 1 times

islamfouda 2 years, 7 months ago

C is right, even if DataViwer is granted on Project level but Dataset is shared with view access to only the country group.

upvoted 2 times

JohnDohertyDoe 2 years, 7 months ago

A is the correct answer. Tested the two scenarios, with `jobUser` permissions it does not allow the user to see a dataset. Whereas with `dataViewer` it has permissions for all the datasets. Note the difference is in the initial permission across the project and not per dataset.

upvoted 7 times

bandegg 2 years, 7 months ago

It's A because in order to query, on needs the jobUser role. dataViewer doesn't grant the ability to actually query the datasets one has been given access to. https://cloud.google.com/bigquery/docs/running-queries#required_permissions

upvoted 1 times

e5019c6 2 years, 7 months ago

I'm siding with C on this one. jobUser role has the bigquery.jobs.create permission, which allow it to load data into BQ, which analyst shouldn't do. Data Viewer has no permissions to add or edit data (It can create a snapshot of the data, extract it or replicate it at most)

upvoted 1 times

whoosh007 2 years, 7 months ago

BigQuery Data Viewer (roles/bigquery.dataViewer) When applied to a table or view, this role provides permissions to: Read data and metadata from the table or view. This role cannot be applied to individual models or routines. When applied to a dataset, this role provides permissions to: Read the dataset's metadata and list tables in the dataset. Read data and metadata from the dataset's tables. When applied at the project or organization level, this role can also enumerate all datasets in the project. Additional roles, however, are necessary to allow the running of jobs. Lowest-level resources where you can grant this role: Table and view BigQuery Job User (roles/bigquery.jobUser) Provides permissions to run jobs, including queries, within the project. Lowest-level resources where you can grant this role: Project Analyst must query data --> BigQuery Data Viewer

upvoted 3 times

steghe 2 years, 10 months ago

A: JobUser to execute queries in general. Data viewer for viewing the country dataset.

upvoted 1 times

TopTalk 2 years, 10 months ago

Lowest-level resources where you can grant this role: dataViewer: Table, View jobUser: Project You don't want to grant access to the entire project, only the dataset which is divided per country. Definitely C. https://cloud.google.com/bigquery/docs/access-control#bigquery.dataViewerCorrect answer is A. Note this: "Share the appropriate dataset with *view access* with each respective analyst country-group". "view access" is the key.

upvoted 2 times

TopTalk 2 years, 10 months ago

Selected Answer: C

Lowest-level resources where you can grant this role: dataViewer: Table, View jobUser: Project

You don't want to grant access to the entire project, only the dataset which is divided per country. Definitely C.
https://cloud.google.com/bigquery/docs/access-control#bigquery.dataViewer

upvoted 2 times

don_v 2 years, 7 months ago

Correct answer is A. Note this: "Share the appropriate dataset with *view access* with each respective analyst country-group". "view access" is the key.

upvoted 1 times

AdityaGupta 2 years, 10 months ago

A. Create a group per country. Add analysts to their respective country-groups. Create a single group 'all_analysts', and add all country-groups as members. Grant the 'all_analysts' group the IAM role of BigQuery jobUser. Share the appropriate dataset with view access with each respective analyst country-group. As all analysts need to execute query, they need JobUser role. They should be restricted to view all datasets (not tables) of respective country.

upvoted 1 times

Question #86 *Topic 1*

You have been engaged by your client to lead the migration of their application infrastructure to GCP. One of their current problems is that the on-premises high performance SAN is requiring frequent and expensive upgrades to keep up with the variety of workloads that are identified as follows: 20 TB of log archives retained for legal reasons; 500 GB of VM boot/data volumes and templates; 500 GB of image thumbnails; 200 GB of customer session state data that allows customers to restart sessions even if off-line for several days.
Which of the following best reflects your recommendations for a cost-effective storage allocation?

A. Local SSD for customer session state data. Lifecycle-managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.

B. Memcache backed by Cloud Datastore for the customer session state data. Lifecycle-managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.

C. Memcache backed by Cloud SQL for customer session state data. Assorted local SSD-backed instances for VM boot/data volumes. Cloud Storage for log archives and thumbnails.

D. Memcache backed by Persistent Disk SSD storage for customer session state data. Assorted local SSD-backed instances for VM boot/data volumes. Cloud Storage for log archives and thumbnails.

Suggested Answer: *D*
Official Answer: D

Community vote distribution

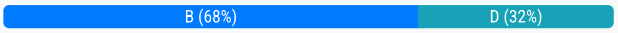

OSNG Highly Voted 5 years, 8 months ago
B is correct. WHY NOT OTHERS. A: is wrong Local SSD in non-persistent therefore cannot be used for session state (as questions also need to save data for users who are offline for several days). C: Again Local SSD cannot be used for boot volume (because its Non-persistent again) and always used for temporary data storage. D: Same reason as C. WHY B? Left with B that's why, but the question is how to store Boot/Data volume on Cloud Storage? - Storing other type of data is easy but most comments were about boot volume. - Boot volume can be stored to Cloud Storage by creating an Custom Image. https://cloud.google.com/compute/docs/images/create-delete-deprecate-private-images#selecting_image_storage_location ---- Upvote if agree for the clarification of others ----
upvoted 114 times

OSNG Highly Voted 5 years, 8 months ago
B is correct. WHY NOT OTHERS. A: is wrong Local SSD in non-persistent therefore cannot be used for session state (as questions also need to save data for users who are offline for several days). C: Again Local SSD cannot be used for boot volume (because its Non-persistent again) and always used for temporary data storage. D: Same reason as C. WHY B? Left with B that's why, but the question is how to store Boot/Data volume on Cloud Storage? - Storing other type of data is easy but most comments were about boot volume. - Boot volume can be stored to Cloud Storage by creating an Custom Image. https://cloud.google.com/compute/docs/images/create-delete-deprecate-private-images#selecting_image_storage_location ---- Upvote if agree for the clarification of others ----How can u use cloud storage for VM boot/data volumnes ? B is wrongAll the options are debatable and have some flaw. But closes answer is B although it has a flaw mentioned above but is still better than other optionsAnswer is D - boot volumes (not boot images) cannot come from Cloud Storage - so B is not the answer.Guys I think its a english error. The last line need to be read carefully Decouple the line after , and Vm boot/data volumes. I think they mean to use vm persistent disks for boot and data volumes. B is the answerYou rejected D with same reason as C. But D is using Persistent disk, not local SSD And not a good idea to keep VM boot data in Cloud storage. As VM needs fast access to boot images to boot upalso Datastore (now Firestore) is aimed for storing stateD states local SSD for VMs. ¨The data that you store on a local SSD persists only until the instance is stopped or deleted¨Yeah, D looks good except the local SSD. Although would we ever want to stop the VMs and not be happy to rebuild? Not sure about B, a move to noSQL is a big, Cloud Storage for the VM boot disks rules it out for sure.
upvoted 114 times

Ishu_awsguy 4 years ago
How can u use cloud storage for VM boot/data volumnes ? B is wrong
upvoted 1 times

Ishu_awsguy 4 years ago
How can u use cloud storage for VM boot/data volumnes ? B is wrongAll the options are debatable and have some flaw. But closes answer is B although it has a flaw mentioned above but is still better than other options
upvoted 1 times

Ishu_awsguy 4 years ago

All the options are debatable and have some flaw. But closes answer is B although it has a flaw mentioned above but is still better than other options
upvoted 3 times

neversaynever 4 years, 7 months ago
Answer is D - boot volumes (not boot images) cannot come from Cloud Storage - so B is not the answer.
upvoted 5 times

Ishu_awsguy 4 years ago
Guys I think its a english error. The last line need to be read carefully Decouple the line after , and Vm boot/data volumes. I think they mean to use vm persistent disks for boot and data volumes. B is the answer
upvoted 2 times

6721sora 3 years, 12 months ago
You rejected D with same reason as C. But D is using Persistent disk, not local SSD And not a good idea to keep VM boot data in Cloud storage. As VM needs fast access to boot images to boot up
upvoted 4 times

6721sora 3 years, 12 months ago
You rejected D with same reason as C. But D is using Persistent disk, not local SSD And not a good idea to keep VM boot data in Cloud storage. As VM needs fast access to boot images to boot upalso Datastore (now Firestore) is aimed for storing stateD states local SSD for VMs. ¨The data that you store on a local SSD persists only until the instance is stopped or deleted¨Yeah, D looks good except the local SSD. Although would we ever want to stop the VMs and not be happy to rebuild? Not sure about B, a move to noSQL is a big, Cloud Storage for the VM boot disks rules it out for sure.
upvoted 4 times

melono 3 years, 10 months ago
also Datastore (now Firestore) is aimed for storing state
upvoted 1 times

melono 3 years, 10 months ago
D states local SSD for VMs. ¨The data that you store on a local SSD persists only until the instance is stopped or deleted¨Yeah, D looks good except the local SSD. Although would we ever want to stop the VMs and not be happy to rebuild? Not sure about B, a move to noSQL is a big, Cloud Storage for the VM boot disks rules it out for sure.
upvoted 1 times

melono 3 years, 10 months ago
D states local SSD for VMs. ¨The data that you store on a local SSD persists only until the instance is stopped or deleted¨
upvoted 1 times

BiddlyBdoyng 3 years, 10 months ago
Yeah, D looks good except the local SSD. Although would we ever want to stop the VMs and not be happy to rebuild? Not sure about B, a move to noSQL is a big, Cloud Storage for the VM boot disks rules it out for sure.
upvoted 1 times

siumk 6 years, 5 months ago
IMHO Answer is B:

Memcache backed by Cloud Datastore https://cloud.google.com/appengine/docs/standard/python/memcache

Compute Engine image can be stored in Cloud Storage https://cloud.google.com/solutions/image-management-best-practices After the complete sequence of bytes from the disk are written to the file, the file is archived using the tar format and then compressed using the GZIP format. You can then upload the resulting *.tar.gz file to Cloud Storage and register it as an image in Compute Engine.
upvoted 17 times

siumk 6 years, 5 months ago
IMHO Answer is B: Memcache backed by Cloud Datastore https://cloud.google.com/appengine/docs/standard/python/memcache Compute Engine image can be stored in Cloud Storage https://cloud.google.com/solutions/image-management-best-practices After the complete sequence of bytes from the disk are written to the file, the file is archived using the tar format and then compressed using the GZIP format. You can then upload the resulting *.tar.gz file to Cloud Storage and register it as an image in Compute Engine.The problem with B is that they are using SAN for data volumes of working VMs, not just to store templates/images. All answers here are quite bad. But I would go with D, as they are talking about several days of saving users' stale session data, which is something that can be accomplished with SSD.@ayzen yes. IS cloud datastore optimized to handle such a data (200GB)
upvoted 17 times

Ayzen 6 years, 3 months ago
The problem with B is that they are using SAN for data volumes of working VMs, not just to store templates/images. All answers here are quite bad. But I would go with D, as they are talking about several days of saving users' stale session data, which is something that can be accomplished with

SSD.

upvoted 10 times

Ayzen 6 years, 3 months ago

The problem with B is that they are using SAN for data volumes of working VMs, not just to store templates/images. All answers here are quite bad. But I would go with D, as they are talking about several days of saving users' stale session data, which is something that can be accomplished with SSD.@ayzen yes. IS cloud datastore optimized to handle such a data (200GB)

upvoted 10 times

Bijesh 5 years, 9 months ago

@ayzen yes. IS cloud datastore optimized to handle such a data (200GB)

upvoted 1 times

JoseCloudEng1994 Most Recent 11 months, 3 weeks ago

Why the Other Options Are Incorrect A. Local SSD for customer session state data... This is incorrect because Local SSD is ephemeral. All data is lost if the VM instance is stopped or terminated. This violates the requirement for session data to be durable for several days. C. ...Assorted local SSD-backed instances for VM boot/data volumes. This is incorrect for the same reason. Local SSD cannot be used for persistent boot volumes. The data would not survive an instance stop/start. This is a critical architectural error. D. ...Assorted local SSD-backed instances for VM boot/data volumes. This suffers from the same fatal flaw as option C. Using ephemeral Local SSD for persistent VM volumes is not a viable solution.

upvoted 1 times

Toothpick 2 years ago

None of these provide an effective method of storing boot/data volumes The correct approach would be to create persistent disk Drives for boot/data volumes directly and go with B or D for the remainder of requirements.

upvoted 2 times

otts 2 years, 7 months ago

B is correct. the question lays emphasis on a cost-effective storage allocation, and persistent disks are costly than B (that rules out D)

upvoted 1 times

e5019c6 2 years, 7 months ago

This is a troublesome question... IMHO, it's D. I was between C & D, but seeing how they stored the customer session data as files, storing it in SQL would require refactoring, and maybe higher latencies. I'm against B because of the latency added by booting & loading data off Cloud Storage, since it adds network latency to the equation. BUT this is assuming the method of using cloud storage is via gcsfuse, that is, using the bucket as a HD. Now if the way of using it is via image of the disk that is loaded when the instance starts, that would be ok. And that's what I expected of option D, that it would load an image of the boot/data volume in it's non-persistent disks. No 'persistent' data would be stored in it, so anything lost when it's shutdown can be ignored.

upvoted 1 times

AdityaGupta 2 years, 10 months ago

D is correct answer because - Memcache backed by Persistent Disk SSD storage for customer session state data. - Persistent disk will ensure that Session data is preserved evern if not used for long time. - Assorted local SSD-backed instances for VM boot/data volumes. Provides faster boot up for VMs. (There is no requirement for persistent storage) - Cloud Storage for log archives and thumbnails. - For low cost and scalable solution.

upvoted 1 times

ductrinh 2 years, 10 months ago

d wrong because of local ssd cannot use for boot image. its temporary and will be cleared if vm was suspended >> so chose bA lot of folks keep saying D is wrong because of "local SSD". It never mentioned "local" in option D. It said "Persistent Disk SSD". https://cloud.google.com/compute/docs/disks I would definitely choose D.My bad, I was focusing on the memcache. I would choose C. The "assorted local ssd" for VM boot/data volumes can't work, as Local SSDs are ephemeral, meaning it's lost if the VM instance is stopped or deleted.I think I'm tired. I meant to write Option B as the correct answer. Memcache backed by Cloud Datastore, etc.

upvoted 2 times

ductrinh 2 years, 10 months ago

d wrong because of local ssd cannot use for boot image. its temporary and will be cleared if vm was suspended >> so chose b

upvoted 2 times

A21325412 2 years, 9 months ago

A lot of folks keep saying D is wrong because of "local SSD". It never mentioned "local" in option D. It said "Persistent Disk SSD". https://cloud.google.com/compute/docs/disks I would definitely choose D.My bad, I was focusing on the memcache. I would choose C. The "assorted local ssd" for VM boot/data volumes can't work, as Local SSDs are ephemeral, meaning it's lost if the VM instance is stopped or deleted.I think I'm tired. I meant to write Option B as the correct answer. Memcache backed by Cloud Datastore, etc.

upvoted 2 times

A21325412 2 years, 9 months ago

A lot of folks keep saying D is wrong because of "local SSD". It never mentioned "local" in option D. It said "Persistent Disk SSD".

https://cloud.google.com/compute/docs/disks

I would definitely choose D.

upvoted 2 times

A21325412 2 years, 9 months ago

My bad, I was focusing on the memcache. I would choose C. The "assorted local ssd" for VM boot/data volumes can't work, as Local SSDs are ephemeral, meaning it's lost if the VM instance is stopped or deleted.I think I'm tired. I meant to write Option B as the correct answer. Memcache backed by Cloud Datastore, etc.

upvoted 1 times

A21325412 2 years, 9 months ago

My bad, I was focusing on the memcache. I would choose C. The "assorted local ssd" for VM boot/data volumes can't work, as Local SSDs are ephemeral, meaning it's lost if the VM instance is stopped or deleted.

upvoted 1 times

A21325412 2 years, 9 months ago

I think I'm tired. I meant to write Option B as the correct answer. Memcache backed by Cloud Datastore, etc.

upvoted 1 times

Murtuza 2 years, 11 months ago

Local SSD means Ephemeral (Temp Disk) so do not confused with PERSISTENT disk (Permanent Disk). The data that you store on a local SSD persists only until the instance is stopped or deleted.

upvoted 1 times

medi01 3 years, 4 months ago

Local SSD cannot be used for neither boot nor data!!! This rules out B&C. Oh, and A too.

upvoted 1 times

Kamaly 3 years, 4 months ago

Cloud Datastore is the right solution to store the session data

upvoted 1 times

feholen210 3 years, 5 months ago

B Seems correct.

upvoted 1 times

ile02 3 years, 5 months ago

D. makes more sense

upvoted 1 times

WAENANY 3 years, 5 months ago

d makes more sense

upvoted 1 times

Deb2293 3 years, 5 months ago

ChatGPT says option D.chatgpt has knowledge till Sept 2021. Don't rely on it bro

upvoted 1 times

Deb2293 3 years, 5 months ago

Selected Answer: D

ChatGPT says option D.

upvoted 1 times

Deb2293 3 years, 5 months ago

chatgpt has knowledge till Sept 2021. Don't rely on it bro

upvoted 5 times

kaleemahmad75 3 years, 7 months ago

My answer is D

upvoted 1 times

Santanu_01 3 years, 7 months ago

I will go with Option D as it is best practice to keep similar data together and seprate OS, Volatile and permanent

upvoted 2 times

Question #87 *Topic 1*

Your web application uses Google Kubernetes Engine to manage several workloads. One workload requires a consistent set of hostnames even after pod scaling and relaunches.
Which feature of Kubernetes should you use to accomplish this?

A. StatefulSets

B. Role-based access control

C. Container environment variables

D. Persistent Volumes

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (100%)

Eroc Highly Voted 5 years, 3 months ago
StatefulSets is a feature of Kubernetes, which the question asks about. Yes, Persistent volumes are required by StatefulSets (https://kubernetes.io/docs/concepts/workloads/controllers/statefulset/). See the Google documentations for mentioning of hostnames (https://cloud.google.com/kubernetes-engine/docs/concepts/statefulset)... Answer A
upvoted 59 times

Eroc Highly Voted 5 years, 3 months ago
StatefulSets is a feature of Kubernetes, which the question asks about. Yes, Persistent volumes are required by StatefulSets (https://kubernetes.io/docs/concepts/workloads/controllers/statefulset/). See the Google documentations for mentioning of hostnames (https://cloud.google.com/kubernetes-engine/docs/concepts/statefulset)... Answer AA is okthank you!It is AA is correct, statefulset
upvoted 59 times

tartar 4 years, 6 months ago
A is ok
upvoted 6 times

OrangeTiger 3 years, 1 month ago
thank you!
upvoted 1 times

nitinz 3 years, 11 months ago
It is A
upvoted 2 times

kumarp6 4 years, 3 months ago
A is correct, statefulset
upvoted 2 times

omermahgoub 2 years, 1 month ago
A. StatefulSets To ensure that a workload in Kubernetes has a consistent set of hostnames even after pod scaling and relaunches, you should use StatefulSets. StatefulSets are a type of controller in Kubernetes that is used to manage stateful applications. They provide a number of features that are specifically designed to support stateful applications, including: Stable, unique network identifiers for each pod in the set Persistent storage that is automatically attached to pods Ordered, graceful deployment and scaling of pods Ordered, graceful deletion and termination of pods By using StatefulSets, you can ensure that your workload has a consistent set of hostnames even if pods are scaled or relaunched, which can be important for applications that rely on stable network identifiers.You are the best.. thanks for all the hard work to explainobviously from chatGPT, but still good to share.
upvoted 20 times

omermahgoub 2 years, 1 month ago
A. StatefulSets

To ensure that a workload in Kubernetes has a consistent set of hostnames even after pod scaling and relaunches, you should use StatefulSets. StatefulSets are a type of controller in Kubernetes that is used to manage stateful applications. They provide a number of features that are

specifically designed to support stateful applications, including:

Stable, unique network identifiers for each pod in the set Persistent storage that is automatically attached to pods Ordered, graceful deployment and scaling of pods Ordered, graceful deletion and termination of pods By using StatefulSets, you can ensure that your workload has a consistent set of hostnames even if pods are scaled or relaunched, which can be important for applications that rely on stable network identifiers.

upvoted 20 times

Tamirm 2 years ago

You are the best.. thanks for all the hard work to explainobviously from chatGPT, but still good to share.

upvoted 3 times

Tamirm 2 years ago

You are the best.. thanks for all the hard work to explain

upvoted 3 times

Wangyu60 Most Recent 1 year, 10 months ago

obviously from chatGPT, but still good to share.

upvoted 1 times

kaleemahmad75 2 years, 1 month ago

A is the answer

upvoted 1 times

megumin 2 years, 3 months ago

A is ok

upvoted 1 times

Deepak31 2 years, 3 months ago

A StatefulSet is the Kubernetes controller used to run the stateful application as containers (Pods) in the Kubernetes cluster. StatefulSets assign a sticky identity—an ordinal number starting from zero—to each Pod instead of assigning random IDs for each replica Pod. A new Pod is created by cloning the previous Pod's data.

upvoted 2 times

AzureDP900 2 years, 4 months ago

this is straight forward question if you know kubernetes concepts. A is rightI do not know Kubernetes

upvoted 1 times

AzureDP900 2 years, 4 months ago

this is straight forward question if you know kubernetes concepts. A is right

upvoted 1 times

zr79 2 years, 4 months ago

I do not know Kubernetes

upvoted 2 times

DrishaS4 2 years, 6 months ago

https://kubernetes.io/docs/concepts/workloads/controllers/statefulset/

upvoted 3 times

mv2000 2 years, 7 months ago

06/30/2022 Exam

upvoted 3 times

haroldbenites 3 years, 2 months ago

Go for A. https://kubernetes.io/docs/concepts/workloads/controllers/statefulset/

upvoted 2 times

vincy2202 3 years, 2 months ago

A is the correct answer

upvoted 2 times

MaxNRG 3 years, 3 months ago

A – StatefulSets StatefulSets are suitable for deploying Kafka, MySQL, Redis, ZooKeeper, and other applications needing unique, persistent identities and stable hostnames. Read more about StatefulSets. https://cloud.google.com/kubernetes-engine/docs/concepts/statefulset C – Container Env Variable, are good if you need to init containers with some static content. E.g. Pod passes to containers its HOSTNAME (where containers are running), namespace and user defined vars. So, env vars is a way for Pod to init containers at start up. But, stable hostnames must be preserved at Pod level via StatefulSets. Defining Env Vars for Container: https://kubernetes.io/docs/tasks/inject-data-application/define-environment-variable-container/

upvoted 6 times

Arjun1983 3 years, 3 months ago

StatefulSets are designed to deploy stateful applications and clustered applications that save data to persistent storage, such as Compute Engine persistent disks. StatefulSets are suitable for deploying Kafka, MySQL, Redis, ZooKeeper, and other applications needing unique, persistent identities and "stable hostnames". Answer is A

upvoted 2 times

victory108 3 years, 9 months ago

A. StatefulSets

upvoted 3 times

un 3 years, 9 months ago

A is correct

upvoted 1 times

Ausias18 3 years, 10 months ago

Answer is A

upvoted 1 times

BhupalS 4 years, 2 months ago

A is the Ans https://kubernetes.io/docs/concepts/workloads/controllers/statefulset/

upvoted 1 times

Chulbul_Pandey 4 years, 2 months ago

StatefulSets for sequencing.. A is correct

upvoted 1 times

Question #88 *Topic 1*

You are using Cloud CDN to deliver static HTTP(S) website content hosted on a Compute Engine instance group. You want to improve the cache hit ratio.
What should you do?

A. Customize the cache keys to omit the protocol from the key.

B. Shorten the expiration time of the cached objects.

C. Make sure the HTTP(S) header ɹ€Cache-Regionɹ€ points to the closest region of your users.

D. Replicate the static content in a Cloud Storage bucket. Point CloudCDN toward a load balancer on that bucket.

Suggested Answer: *A*
Reference: https://cloud.google.com/cdn/docs/best-practices#using_custom_cache_keys_to_improve_cache_hit_ratio

Community vote distribution
A (70%) D (30%)

shandy Highly Voted 6 years, 8 months ago
Option A is Correct. https://cloud.google.com/cdn/docs/caching#cache-keysA is okYes, A is correctA, both http and https will use same key.https://cloud.google.com/cdn/docs/best-practices#cache-hit-ratio
upvoted 24 times

shandy Highly Voted 6 years, 8 months ago
Option A is Correct. https://cloud.google.com/cdn/docs/caching#cache-keys
upvoted 24 times

tartar 6 years ago
A is ok
upvoted 7 times

kumarp6 5 years, 9 months ago
Yes, A is correct
upvoted 2 times

nitinz 5 years, 5 months ago
A, both http and https will use same key.
upvoted 3 times

MestreCholas 3 years, 5 months ago
https://cloud.google.com/cdn/docs/best-practices#cache-hit-ratio
upvoted 7 times

gfhbox0083 6 years, 2 months ago
A, for sure. By default, Cloud CDN uses the complete request URL to build the cache key. For performance and scalability, it's important to optimize cache hit ratio. To help optimize your cache hit ratio, you can use custom cache keys .....
upvoted 10 times

Ekramy_Elnaggar 1 year, 9 months ago
Cache Keys and Protocols: Cloud CDN uses cache keys to identify and store content in its cache. By default, the protocol (HTTP or HTTPS) is included in the cache key. This means that the same content served over HTTP and HTTPS will be cached separately, reducing the cache hit ratio. Omitting the Protocol: Customizing the cache keys to omit the protocol allows Cloud CDN to treat HTTP and HTTPS requests for the same content as identical. This increases the chance of a cache hit, as the CDN can serve the cached content regardless of the protocol used in the request. Improved Cache Hit Ratio: By consolidating the cache entries for HTTP and HTTPS versions of the content, you effectively increase the cache hit ratio. This leads to better performance, reduced latency, and lower costs.
upvoted 5 times

nareshthumma 1 year, 9 months ago
Answer is DIncorrect. While Cloud Storage is often an excellent origin for static content, simply migrating the content from a Compute Engine instance group to a Cloud Storage bucket does not by itself guarantee an improved cache hit ratio. The hit ratio depends on how the content is served and cached, not just where it's stored. The core issue remains how the cache keys are configured.
upvoted 1 times

nareshthumma 1 year, 9 months ago

Answer is D

upvoted 1 times

JoseCloudEng1994 Most Recent 11 months, 3 weeks ago

Incorrect. While Cloud Storage is often an excellent origin for static content, simply migrating the content from a Compute Engine instance group to a Cloud Storage bucket does not by itself guarantee an improved cache hit ratio. The hit ratio depends on how the content is served and cached, not just where it's stored. The core issue remains how the cache keys are configured.

upvoted 1 times

Prajjwal199831 2 years, 5 months ago

A is right

upvoted 2 times

JB28 2 years, 6 months ago

option D

upvoted 2 times

bargou 2 years, 6 months ago

i'm with D Option

upvoted 2 times

AdityaGupta 2 years, 10 months ago

https://cloud.google.com/cdn/docs/best-practices#using_custom_cache_keys_to_improve_cache_hit_ratio A is correct option, use custom keys.

upvoted 2 times

Gregwaw 2 years, 10 months ago

A and D could be OK but there is no information in the question that both http and https protocols are used (it is in fact not probable to use both protocols). Anwer A would be only valid when both http and https are in use.Incorrect, respectfully. "HTTP(S)" (the parenthesis) implies that HTTP requests can be made using *either* HTTP or HTTPS, especially if TLS is not mandatory. If a site returns a 200 OK for http://www.example.com/picture/of/a/cat.jpg and httpS://www.example.come/picture/of/a/cat.jpg, then you should exclude the protocol in the cache key to increase the cache-hit ratio. If you don't, the CDN will treat both URLs as different objects.

upvoted 1 times

Gregwaw 2 years, 10 months ago

Selected Answer: D

A and D could be OK but there is no information in the question that both http and https protocols are used (it is in fact not probable to use both protocols). Anwer A would be only valid when both http and https are in use.

upvoted 1 times

MikeH20 2 years, 8 months ago

Incorrect, respectfully. "HTTP(S)" (the parenthesis) implies that HTTP requests can be made using *either* HTTP or HTTPS, especially if TLS is not mandatory. If a site returns a 200 OK for http://www.example.com/picture/of/a/cat.jpg and httpS://www.example.come/picture/of/a/cat.jpg, then you should exclude the protocol in the cache key to increase the cache-hit ratio. If you don't, the CDN will treat both URLs as different objects.

upvoted 2 times

heretolearnazure 2 years, 12 months ago

A sounds correct

upvoted 1 times

daidaidai 2 years, 12 months ago

Replicating static content in a Cloud Storage bucket and pointing CloudCDN toward a load balancer on that bucket may improve the distribution of content but is not directly related to improving the cache hit ratio based on the customizing of cache keys.

upvoted 1 times

VarunGo 3 years, 3 months ago

D This option is the best because Cloud Storage has built-in caching and can serve content faster than Compute Engine instances. It also allows for better scalability and availability. By pointing Cloud CDN towards a load balancer on the Cloud Storage bucket, the cache hit ratio can be improved as the content will be served directly from the cache without needing to access the Compute Engine instances. Option A (Customize the cache keys to omit the protocol from the key) may not be effective in improving the cache hit ratio as it only removes the protocol from the cache key and does not address the underlying issue of slow content delivery.

upvoted 2 times

Deb2293 3 years, 5 months ago

Customizing the cache keys by omitting the protocol from the key (option A) can be a valid approach to improve the cache hit ratio for CDN delivered Compute Engine, but it may not be the most effective solution for all cases. Customizing the cache keys can improve the cache hit ratio by reducing

the number of cache misses caused by variations in the request URL, headers, and parameters. However, customizing the cache keys requires careful consideration of the caching policies, traffic patterns, and content types

upvoted 1 times

omermahgoub 3 years, 8 months ago

A. Customize the cache keys to omit the protocol from the key. To improve the cache hit ratio with Cloud CDN, you should customize the cache keys to omit the protocol (e.g. HTTP or HTTPS) from the key. This will allow Cloud CDN to cache the same content under both HTTP and HTTPS, which can help to improve the hit ratio by allowing Cloud CDN to serve content from cache more frequently. To customize the cache keys, you can use the --key-include-protocol flag when enabling Cloud CDN for your Compute Engine instance group or load balancer. Setting this flag to false will cause Cloud CDN to omit the protocol from the cache key. Other options, such as shortening the expiration time of cached objects or replicating content in Cloud Storage, may also help to improve the cache hit ratio, but customizing the cache keys to omit the protocol is likely to have the greatest impact.

upvoted 8 times

AzureDP900 3 years, 10 months ago

https://cloud.google.com/cdn/docs/best-practices#cache-hit-ratio. A is rightEach cache entry in a Cloud CDN cache is identified by a cache key. When a request comes into the cache, the cache converts the URI of the request into a cache key, and then compares it with keys of cached entries. If it finds a match, the cache returns the object associated with that key.

upvoted 3 times

AzureDP900 3 years, 10 months ago

https://cloud.google.com/cdn/docs/best-practices#cache-hit-ratio. A is right

upvoted 3 times

AzureDP900 3 years, 10 months ago

Each cache entry in a Cloud CDN cache is identified by a cache key. When a request comes into the cache, the cache converts the URI of the request into a cache key, and then compares it with keys of cached entries. If it finds a match, the cache returns the object associated with that key.

upvoted 2 times

aut0pil0t 3 years, 11 months ago

Use case: "A logo needs to be cached whether displayed through HTTP or HTTPS. When you customize the cache keys for the backend service that holds the logo, clear the Protocol checkbox so that requests through HTTP and HTTPS count as matches for the logo's cache entry."
https://cloud.google.com/cdn/docs/best-practices#using_custom_cache_keys_to_improve_cache_hit_ratio

upvoted 2 times

DrishaS4 4 years ago

https://cloud.google.com/cdn/docs/best-practices#using_custom_cache_keys_to_improve_cache_hit_ratio

upvoted 1 times

Question #89 *Topic 1*

Your architecture calls for the centralized collection of all admin activity and VM system logs within your project.
How should you collect these logs from both VMs and services?

A. All admin and VM system logs are automatically collected by Stackdriver.

B. Stackdriver automatically collects admin activity logs for most services. The Stackdriver Logging agent must be installed on each instance to collect system logs.

C. Launch a custom syslogd compute instance and configure your GCP project and VMs to forward all logs to it.

D. Install the Stackdriver Logging agent on a single compute instance and let it collect all audit and access logs for your environment.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
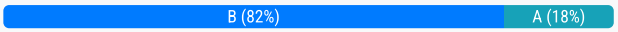

MeasService Highly Voted 6 years, 4 months ago
Does not agree with D. B is the nearest answer I feel !
upvoted 43 times

MeasService Highly Voted 6 years, 4 months ago
Does not agree with D. B is the nearest answer I feel !B is correct, D is SPOF ...It is B, all rest are BSAgree.B is ok
upvoted 43 times

kumarp6 5 years, 3 months ago
B is correct, D is SPOF ...
upvoted 2 times

nitinz 4 years, 11 months ago
It is B, all rest are BS
upvoted 2 times

KouShikyou 6 years, 4 months ago
Agree.
upvoted 9 times

KouShikyou 6 years, 4 months ago
Agree.B is ok
upvoted 9 times

tartar 5 years, 6 months ago
B is ok
upvoted 12 times

shandy 6 years, 2 months ago
Admin and event logs are configured by default. VM System logs require a logging agent to be configured. So A is not valid. Answer is B
upvoted 21 times

JaimeMS Most Recent 1 year, 8 months ago
Admin and event logs are configured by default. VM System logs require a logging agent to be configured. Answer is B
upvoted 1 times

hitmax87 1 year, 9 months ago
Stackdriver Logging agent doesn't need for audit logs
upvoted 1 times

convers39 2 years, 1 month ago
B is correct Now it is recommended to use OpsAgent as a replacement. Although you can create a VM instance with OpsAgent automatically enabled, which makes it look like 'the logging is automatically enabled', under the hood you need to install the agent on the instance.
https://cloud.google.com/stackdriver/docs/solutions/agents/ops-agent/install-agent-vm-creation
upvoted 4 times

stefanop 2 years, 2 months ago

A: Stackdriver already collects admin logs and GCE logs.

upvoted 2 times

spuyol 2 years, 2 months ago

Answer is A you don't need to install the Stackdriver Logging agent or any other agents in order to collect admin or system logs from compute engine instances (VM): https://cloud.google.com/compute/docs/logging/activity-logs https://cloud.google.com/compute/docs/logging/audit-logging

upvoted 2 times

squishy_fishy 2 years, 3 months ago

Answer is B, but it would be more accurate if the answer mentioned install the Ops Agent on the VMs.
https://cloud.google.com/logging/docs/agent/ops-agent

upvoted 1 times

jlambdan 2 years, 11 months ago

answer is B as per this tutorial step: https://cloud.google.com/logging/docs/logging-gce-quickstart#install-agent

upvoted 1 times

omermahgoub 3 years, 1 month ago

Stackdriver does not require the Stackdriver Logging agent to be installed in order to collect system logs. Stackdriver is a cloud monitoring and logging platform that is integrated with Google Cloud Platform (GCP) and is designed to collect, monitor, and troubleshoot logs from your GCP resources. By default, Stackdriver automatically collects admin activity logs for most GCP services, as well as VM system logs. This means that you don't need to install the Stackdriver Logging agent or any other agents in order to collect these logs - they are automatically collected and centralized by Stackdriver. However, if you want to collect logs from other sources that are not automatically collected by Stackdriver (e.g. logs from applications running on your VMs, logs from on-premises systems, etc.), you can use the Stackdriver Logging agent to forward these logs to Stackdriver. The agent is a lightweight daemon that runs on your VMs or other hosts, and it can be used to collect logs from various sources and forward them to Stackdriver for centralized storage and analysis.Answer is A In Google Cloud Platform (GCP), you can use Stackdriver to collect and centralize all admin activity and VM system logs within your project. Stackdriver is a powerful cloud monitoring and logging platform that is integrated with GCP, and it provides a number of features that are specifically designed to help you collect, monitor, and troubleshoot logs from your GCP resources. One of the key features of Stackdriver is that it automatically collects admin activity logs for most GCP services, as well as VM system logs. This means that you don't need to install any agents or configure any additional components to collect these logs - they are automatically collected and centralized by Stackdriver. To view and analyze your logs in Stackdriver, you can use the Stackdriver Logs Viewer, which provides a powerful interface for searching, filtering, and aggregating your logs. You can also use the Stackdriv

upvoted 6 times

omermahgoub 3 years, 1 month ago

Stackdriver does not require the Stackdriver Logging agent to be installed in order to collect system logs. Stackdriver is a cloud monitoring and logging platform that is integrated with Google Cloud Platform (GCP) and is designed to collect, monitor, and troubleshoot logs from your GCP resources. By default, Stackdriver automatically collects admin activity logs for most GCP services, as well as VM system logs. This means that you don't need to install the Stackdriver Logging agent or any other agents in order to collect these logs - they are automatically collected and centralized by Stackdriver.

However, if you want to collect logs from other sources that are not automatically collected by Stackdriver (e.g. logs from applications running on your VMs, logs from on-premises systems, etc.), you can use the Stackdriver Logging agent to forward these logs to Stackdriver. The agent is a lightweight daemon that runs on your VMs or other hosts, and it can be used to collect logs from various sources and forward them to Stackdriver for centralized storage and analysis.

upvoted 6 times

omermahgoub 3 years, 1 month ago

Answer is A In Google Cloud Platform (GCP), you can use Stackdriver to collect and centralize all admin activity and VM system logs within your project. Stackdriver is a powerful cloud monitoring and logging platform that is integrated with GCP, and it provides a number of features that are specifically designed to help you collect, monitor, and troubleshoot logs from your GCP resources. One of the key features of Stackdriver is that it automatically collects admin activity logs for most GCP services, as well as VM system logs. This means that you don't need to install any agents or configure any additional components to collect these logs - they are automatically collected and centralized by Stackdriver. To view and analyze your logs in Stackdriver, you can use the Stackdriver Logs Viewer, which provides a powerful interface for searching, filtering, and aggregating your logs. You can also use the Stackdriver Logs API to programmatically access your logs, or use the Stackdriver Logging agent to forward your logs to other log management or analysis tools.

upvoted 6 times

megumin 3 years, 3 months ago

B is ok

upvoted 1 times

DrishaS4 3 years, 6 months ago

Selected Answer: B

https://cloud.google.com/logging/docs/agent/logging/installation#before_you_begin

upvoted 5 times

DrishaS4 3 years, 6 months ago

https://cloud.google.com/logging/docs/agent/logging/installation#before_you_beginThank you for sharing the link, B is right

upvoted 5 times

AzureDP900 3 years, 4 months ago

Thank you for sharing the link, B is right

upvoted 1 times

panqueca 3 years, 9 months ago

it's B

upvoted 2 times

haroldbenites 4 years, 2 months ago

Go for B.

upvoted 2 times

vincy2202 4 years, 2 months ago

B is the correct answer.

upvoted 2 times

MaxNRG 4 years, 3 months ago

B – stackdriver automatically collects admin activity logs for most services. Stackdriver Logging Agenct must be installed on each instance to collect system logs. Read more about Logging Agent. https://cloud.google.com/logging/docs/agent/ Logging agent streams logs from 3rd party apps and systems SW (syslog on Linux) to Logging. It is best practice to run the Logging agent on all your VM instances. It runs on Linux and Windows. Cloud Audit Logs says that Admin Activity audit logs are always enabled.

upvoted 3 times

unnikrisb 4 years, 4 months ago

Agree with B

upvoted 1 times

Question #90 *Topic 1*

You have an App Engine application that needs to be updated. You want to test the update with production traffic before replacing the current application version.
What should you do?

A. Deploy the update using the Instance Group Updater to create a partial rollout, which allows for canary testing.

B. Deploy the update as a new version in the App Engine application, and split traffic between the new and current versions.

C. Deploy the update in a new VPC, and use Google's global HTTP load balancing to split traffic between the update and current applications.

D. Deploy the update as a new App Engine application, and use Google's global HTTP load balancing to split traffic between the new and current applications.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (100%)

KouShikyou Highly Voted 5 years, 4 months ago
I think B is correct. Because GAE supports service version control and A/B test. Is my understanding correct?Yes, B is correctOnly B works.
upvoted 61 times

KouShikyou Highly Voted 5 years, 4 months ago
I think B is correct. Because GAE supports service version control and A/B test. Is my understanding correct?
upvoted 61 times

kumarp6 4 years, 3 months ago
Yes, B is correct
upvoted 5 times

nitinz 3 years, 11 months ago
Only B works.
upvoted 4 times

ADVIT 5 years ago
Only one App Engine application can be created per Project. So it's B.
upvoted 15 times

omermahgoub Most Recent 2 years, 1 month ago
B. Deploy the update as a new version in the App Engine application, and split traffic between the new and current versions.

To test an update to an App Engine application with production traffic before replacing the current version, you can deploy the update as a new version in the App Engine application and split traffic between the new and current versions. This is known as a "blue-green" deployment, and it allows you to test the new version with a portion of production traffic while the current version is still serving the remainder of traffic.

To split traffic between the new and current versions, you can use the App Engine traffic splitting feature. This feature allows you to specify the percentage of traffic that should be sent to each version, and it can be used to gradually ramp up traffic to the new version over time. This allows you to test the new version with a small portion of traffic initially, and gradually increase the traffic as you become more confident in the update.
upvoted 11 times

omermahgoub 2 years, 1 month ago
B. Deploy the update as a new version in the App Engine application, and split traffic between the new and current versions. To test an update to an App Engine application with production traffic before replacing the current version, you can deploy the update as a new version in the App Engine application and split traffic between the new and current versions. This is known as a "blue-green" deployment, and it allows you to test the new version with a portion of production traffic while the current version is still serving the remainder of traffic. To split traffic between the new and current versions, you can use the App Engine traffic splitting feature. This feature allows you to specify the percentage of traffic that should be sent to each version, and it can be used to gradually ramp up traffic to the new version over time. This allows you to test the new version with a small portion of traffic initially, and gradually increase the traffic as you become more confident in the update.Other options, such as deploying the update in a new

VPC or as a new App Engine application, are not recommended for testing updates with production traffic, as they can be more complex and may require additional steps to set up.

upvoted 11 times

omermahgoub 2 years, 1 month ago

Other options, such as deploying the update in a new VPC or as a new App Engine application, are not recommended for testing updates with production traffic, as they can be more complex and may require additional steps to set up.

upvoted 3 times

TonytheTiger 2 years, 2 months ago

Answer B : You can use traffic splitting to specify a percentage distribution of traffic across two or more of the versions within a service. Splitting traffic allows you to conduct A/B testing between your versions and provides control over the pace when rolling out features. Traffic splitting is applied to URLs that do not explicitly target a version. For example, the following URLs split traffic because they target all the available versions within the specified service: https://cloud.google.com/appengine/docs/standard/splitting-traffic

upvoted 3 times

megumin 2 years, 3 months ago

B is ok

upvoted 2 times

AzureDP900 2 years, 4 months ago

B is right , Option D is just to confuse you. Deploy the update as a new version in the App Engine application, and split traffic between the new and current versions.

upvoted 2 times

DrishaS4 2 years, 6 months ago

Versioning is supported in App Engine.

upvoted 2 times

ghadxx 3 years ago

Versioning is supported in App Engine.

upvoted 2 times

haroldbenites 3 years, 2 months ago

Go for D, The option B don´t say with wich service will split the traffic. The option D gives more datail and makes sense.

upvoted 1 times

haroldbenites 3 years, 2 months ago

Go for D, The option B don´t say with wich service will split the traffic. The option D gives more datail and makes sense.No mate, only one app engine per project can be deployed, you can have multiple version on the same app tho. D is to confuse you. B is the only feasible answer in here.

upvoted 1 times

ale_brd_111 2 years, 2 months ago

No mate, only one app engine per project can be deployed, you can have multiple version on the same app tho. D is to confuse you. B is the only feasible answer in here.

upvoted 4 times

vincy2202 3 years, 2 months ago

B is the correct answer

upvoted 1 times

robotgeek 3 years, 3 months ago

A is not because "Instance Group Updater " is only for Computer Engine MIG

upvoted 2 times

MaxNRG 3 years, 3 months ago

B – Deploy the update as a new version in AppEngine app, and split traffic between the new and current versions. Traffic Splitting is feature of AppEngine for A/B testing. https://cloud.google.com/appengine/docs/standard/python/splitting-traffic

upvoted 6 times

unnikrisb 3 years, 4 months ago

B is correct... Canary Testing -> Traffic Splitting

upvoted 1 times

victory108 3 years, 9 months ago

B. Deploy the update as a new version in the App Engine application, and split traffic between the new and current versions

upvoted 3 times

un 3 years, 9 months ago

B is the answer

upvoted 1 times

getzsagar 3 years, 10 months ago

Answer - B Configure how much traffic the version that you just deployed should receive. By default, the initial version that you deploy to your App Engine application is automatically configured to receive 100% of traffic. However, all subsequent versions that you deploy to that same App Engine application must be manually configured, otherwise they receive no traffic. For details about how to configure traffic for your versions, see Migrating and Splitting Traffic. https://cloud.google.com/appengine/docs/admin-api/migrating-splitting-traffic

upvoted 3 times

lynx256 3 years, 10 months ago

B is ok

upvoted 1 times

Question #91 *Topic 1*

All Compute Engine instances in your VPC should be able to connect to an Active Directory server on specific ports. Any other traffic emerging from your instances is not allowed. You want to enforce this using VPC firewall rules.
How should you configure the firewall rules?

A. Create an egress rule with priority 1000 to deny all traffic for all instances. Create another egress rule with priority 100 to allow the Active Directory traffic for all instances.

B. Create an egress rule with priority 100 to deny all traffic for all instances. Create another egress rule with priority 1000 to allow the Active Directory traffic for all instances.

C. Create an egress rule with priority 1000 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 100 to block all traffic for all instances.

D. Create an egress rule with priority 100 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 1000 to block all traffic for all instances.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (100%)

wk Highly Voted 6 years, 4 months ago
Should be A, there is no implied deny egress but only implied allow egress

https://cloud.google.com/vpc/docs/firewalls#default_firewall_rules

Every VPC network has two implied firewall rules. These rules exist, but are not shown in the Cloud Console:

The implied allow egress rule: An egress rule whose action is allow, destination is 0.0.0.0/0, and priority is the lowest possible (65535) lets any instance send traffic to any destination, except for traffic blocked by GCP. Outbound access may be restricted by a higher priority firewall rule. Internet access is allowed if no other firewall rules deny outbound traffic and if the instance has an external IP address or uses a NAT instance. Refer to Internet access requirements for more details.

The implied deny ingress rule: An ingress rule whose action is deny, source is 0.0.0.0/0, and priority is the lowest possible (65535) protects all instances by blocking incoming traffic to them. Incoming access may be allowed by a higher priority rule. Note that the default network includes some additional rules that override this one, allowing certain types of incoming traffic.
upvoted 94 times

wk Highly Voted 6 years, 4 months ago
Should be A, there is no implied deny egress but only implied allow egress https://cloud.google.com/vpc/docs/firewalls#default_firewall_rules Every VPC network has two implied firewall rules. These rules exist, but are not shown in the Cloud Console: The implied allow egress rule: An egress rule whose action is allow, destination is 0.0.0.0/0, and priority is the lowest possible (65535) lets any instance send traffic to any destination, except for traffic blocked by GCP. Outbound access may be restricted by a higher priority firewall rule. Internet access is allowed if no other firewall rules deny outbound traffic and if the instance has an external IP address or uses a NAT instance. Refer to Internet access requirements for more details. The implied deny ingress rule: An ingress rule whose action is deny, source is 0.0.0.0/0, and priority is the lowest possible (65535) protects all instances by blocking incoming traffic to them. Incoming access may be allowed by a higher priority rule. Note that the default network includes some additional rules that override this one, allowing certain types of incoming traffic.B is correct...from a book: "Firewall rules control network traffic by blocking or allowing traffic into (ingress) or out of (egress) a network. Two implied firewall rules are defined with VPCs: one blocks all incoming traffic, and the other allows all outgoing traffic. You can change this behavior Virtual Private Clouds 115 116 Chapter 6 ■ Designing Networks by defining firewall rules with higher priority. Firewall rules have a priority specified by an integer from 0 to 65535, with 0 being the highest priority and 65535 being the lowest." so this confirms AGood summary. To the point!It is A, rest all do not make sense. If you think of any other option then go back and read about firewalls. Seriously you are not ready for this exam.thank you
upvoted 94 times

kumarp6 5 years, 3 months ago

B is correct...

upvoted 2 times

p4 5 years, 3 months ago

from a book: "Firewall rules control network traffic by blocking or allowing traffic into (ingress) or out of (egress) a network. Two implied firewall rules are defined with VPCs: one blocks all incoming traffic, and the other allows all outgoing traffic. You can change this behavior Virtual Private Clouds 115 116 Chapter 6 ■ Designing Networks by defining firewall rules with higher priority. Firewall rules have a priority specified by an integer from 0 to 65535, with 0 being the highest priority and 65535 being the lowest."

so this confirms A

upvoted 13 times

p4 5 years, 3 months ago

from a book: "Firewall rules control network traffic by blocking or allowing traffic into (ingress) or out of (egress) a network. Two implied firewall rules are defined with VPCs: one blocks all incoming traffic, and the other allows all outgoing traffic. You can change this behavior Virtual Private Clouds 115 116 Chapter 6 ■ Designing Networks by defining firewall rules with higher priority. Firewall rules have a priority specified by an integer from 0 to 65535, with 0 being the highest priority and 65535 being the lowest." so this confirms AGood summary. To the point!

upvoted 13 times

SSS987 2 years, 1 month ago

Good summary. To the point!

upvoted 1 times

nitinz 4 years, 11 months ago

It is A, rest all do not make sense. If you think of any other option then go back and read about firewalls. Seriously you are not ready for this exam.

upvoted 2 times

nitinz 4 years, 11 months ago

It is A, rest all do not make sense. If you think of any other option then go back and read about firewalls. Seriously you are not ready for this exam.thank you

upvoted 2 times

zr79 3 years, 4 months ago

thank you

upvoted 1 times

MeasService 6 years, 4 months ago

Agree Correct is A. There is no implied deny egress only deny ingress ruleAgree with A . only Implied allow egress rule (or) Implied deny ingress rule. There is No "Implied deny egress rule" which rules out C & D

upvoted 10 times

MeasService 6 years, 4 months ago

Agree Correct is A. There is no implied deny egress only deny ingress rule

upvoted 10 times

MyPractice 6 years, 1 month ago

Agree with A . only Implied allow egress rule (or) Implied deny ingress rule. There is No "Implied deny egress rule" which rules out C & D

upvoted 3 times

Pime13 Most Recent 6 months ago

Selected Answer: A

Priority-Based Evaluation: Google Cloud VPC firewall rules are evaluated based on their priority number, from 0 to 65535. A lower numerical value signifies a higher priority. The first rule that matches the traffic is the one that gets applied.

Overriding Implied Rules: Every VPC network has an implied allow egress rule with the lowest possible priority (65535) that permits all outbound traffic by default. To block "all other traffic," you must create a custom rule with a higher priority (a lower number like 1000) to override this default behavior.

upvoted 1 times

Pime13 6 months ago

Priority-Based Evaluation: Google Cloud VPC firewall rules are evaluated based on their priority number, from 0 to 65535. A lower numerical value signifies a higher priority. The first rule that matches the traffic is the one that gets applied. Overriding Implied Rules: Every VPC network has an implied allow egress rule with the lowest possible priority (65535) that permits all outbound traffic by default. To block "all other traffic," you must create a custom rule with a higher priority (a lower number like 1000) to override this default behavior.Selective Allowance (Least Privilege): To allow the specific Active Directory traffic, you need a rule with an even higher priority (a lower number like 100) than your "deny all" rule. This ensures that

when an instance attempts to reach the AD server, it hits the "allow" rule (priority 100) first. If the traffic is destined anywhere else, it will miss the first rule and be caught by the "deny all" rule (priority 1000).

upvoted 1 times

Pime13 6 months ago

Selective Allowance (Least Privilege): To allow the specific Active Directory traffic, you need a rule with an even higher priority (a lower number like 100) than your "deny all" rule. This ensures that when an instance attempts to reach the AD server, it hits the "allow" rule (priority 100) first. If the traffic is destined anywhere else, it will miss the first rule and be caught by the "deny all" rule (priority 1000).

upvoted 1 times

ManishKS 2 years, 6 months ago

B. Create an egress rule with priority 100 to deny all traffic for all instances. Create another egress rule with priority 1000 to allow the Active Directory traffic for all instances. This option creates a deny all rule with a lower priority and an allow rule with a higher priority. This option will work as intended, as the Active Directory traffic will be allowed and all other outbound traffic will be blocked.

upvoted 2 times

Emmarof 2 years, 11 months ago

The answer to this question is A. Explanation: To enforce the requirement that all Compute Engine instances in your VPC should be able to connect to an Active Directory server on specific ports while blocking any other traffic emerging from instances, the following two egress rules should be created: Create an egress rule with priority 1000 to deny all traffic for all instances. Create another egress rule with priority 100 to allow the Active Directory traffic for all instances. In this configuration, the rule that allows the AD traffic has a lower priority number than the rule that denies all other traffic. Therefore, this rule should be evaluated first.

upvoted 2 times

Deb2293 2 years, 11 months ago

It should be A. It cannot be D as The Implied allow egress rule, with its action of “allow”, allows all traffic out to the 0.0. 0.0/0 destination, which basically means everywhere. The priority of the implied allow egress rule is the lowest possible, 65535. The implied deny ingress rule, with an action of “deny”, blocks all incoming connections.

upvoted 1 times

8d31d36 3 years ago

The correct answer is B. To enforce that all Compute Engine instances in a VPC can connect to an Active Directory server on specific ports while blocking any other traffic, you should create an egress rule with a high priority (lower numerical value) to deny all traffic from all instances, and another egress rule with a lower priority (higher numerical value) to allow traffic to the Active Directory server on the specific ports. Option B creates the necessary egress rules in the correct order: a deny-all rule with a high priority (100), followed by an allow rule for the Active Directory traffic with a lower priority (1000). This way, traffic to the Active Directory server is allowed, but all other traffic is denied.

upvoted 1 times

megumin 3 years, 3 months ago

A is ok

upvoted 1 times

AzureDP900 3 years, 4 months ago

It is pretty straight forward question, It this case priority low should be allow and high priority rules deny all requests. A is right

upvoted 2 times

DrishaS4 3 years, 6 months ago

https://cloud.google.com/vpc/docs/firewalls#priority_order_for_firewall_rules

upvoted 1 times

mv2000 3 years, 7 months ago

06/30/2022 Exam question.

upvoted 6 times

mv2000 3 years, 7 months ago

06/30/2022 Exam question.Oh, really? I got this question on my exam 2 years ago, I did not expect to repeat this kind of questions in the current exam.

upvoted 6 times

moiradavis 3 years, 7 months ago

Oh, really? I got this question on my exam 2 years ago, I did not expect to repeat this kind of questions in the current exam.

upvoted 1 times

Baumster 3 years, 11 months ago

OT: why is there no way to mark questions for review/repeat later on?

upvoted 1 times

haroldbenites 4 years, 2 months ago

Go for A. While the priority is higher, the egress rule is more restricted. While the priority is higher, the ingress rule is more free.

upvoted 1 times

vincy2202 4 years, 2 months ago

A is correct answer

upvoted 1 times

vchrist 4 years, 2 months ago

to understand rules priority: https://cloud.google.com/vpc/docs/firewalls#priority_order_for_firewall_rules

upvoted 1 times

nqthien041292 4 years, 2 months ago

Vote A

upvoted 1 times

MaxNRG 4 years, 3 months ago

A – create an egress rule with priority 1000 to deny all traffic for all instances. Create another egress rule with priority 100 to allow the Active Directory traffic for all instances. Default Firewall rules (aka implied rules) are following: 1) Egress traffic is allowed to all IP/ports. 2) Ingress traffic is disabled completely. Both these rules have lowest priority (65535) and cannot be removed.
https://cloud.google.com/vpc/docs/firewalls#default_firewall_rules

upvoted 2 times

Question #92 *Topic 1*

Your customer runs a web service used by e-commerce sites to offer product recommendations to users. The company has begun experimenting with a machine learning model on Google Cloud Platform to improve the quality of results.
What should the customer do to improve their model's results over time?

A. Export Cloud Machine Learning Engine performance metrics from Stackdriver to BigQuery, to be used to analyze the efficiency of the model.

B. Build a roadmap to move the machine learning model training from Cloud GPUs to Cloud TPUs, which offer better results.

C. Monitor Compute Engine announcements for availability of newer CPU architectures, and deploy the model to them as soon as they are available for additional performance.

D. Save a history of recommendations and results of the recommendations in BigQuery, to be used as training data.

Suggested Answer: *D*
Official Answer: D

Community vote distribution
D (100%)

ghadxx Highly Voted 3 years ago
Selected Answer: D
Model performance is generally based on the volume of its training data input. The more the data, the better the model.
upvoted 21 times

ghadxx Highly Voted 3 years ago
Model performance is generally based on the volume of its training data input. The more the data, the better the model.I agree with you, D is rightYes, correctly said..This is actually a question for Data Engineer role
upvoted 21 times

AzureDP900 2 years, 4 months ago
I agree with you, D is right
upvoted 1 times

Sur_Nikki 1 year, 9 months ago
Yes, correctly said..This is actually a question for Data Engineer role
upvoted 1 times

sgofficial 2 years, 6 months ago
A,B,C is defining about the performance of ML but not the result....only the training data will give good ML result/predictions
upvoted 6 times

JoseCloudEng1994 Most Recent 11 months, 3 weeks ago
More data = more training data = better model (most of the time)
upvoted 1 times

Deb2293 1 year, 11 months ago
Best answer is D. Other 3 makes no sense
upvoted 1 times

PST21 2 years ago
Need to improve the model results and not performance .. hence D
upvoted 1 times

surajkrishnamurthy 2 years, 2 months ago
Answer is D
upvoted 1 times

DrishaS4 2 years, 6 months ago
Model performance is generally based on the volume of its training data input. The more the data, the better the model.
upvoted 3 times

sivre 2 years, 10 months ago
The following insights and recommendations can be exported (to bigquery): IAM recommender VM machine type recommender Managed instance group machine type recommender Idle PD recommender Idle VM recommender Cloud SQL overprovisioned instance recommender Cloud SQL idle

instance recommender Unattended project recommender Cloud Run Service Identity recommender https://cloud.google.com/recommender/docs/bq-export/export-recommendations-to-bq None of this is correlated with Machine Learning, how can be D? looks more A the answerwhat we will do with metrics , it won't improve our Machine learning model , D is the closest answer , also it didn't say export it said Save , which could be manually moving the data to BQ

upvoted 2 times

sivre 2 years, 10 months ago

The following insights and recommendations can be exported (to bigquery): IAM recommender VM machine type recommender Managed instance group machine type recommender Idle PD recommender Idle VM recommender Cloud SQL overprovisioned instance recommender Cloud SQL idle instance recommender Unattended project recommender Cloud Run Service Identity recommender https://cloud.google.com/recommender/docs/bq-export/export-recommendations-to-bq

None of this is correlated with Machine Learning, how can be D? looks more A the answer

upvoted 2 times

kimharsh 2 years, 10 months ago

what we will do with metrics , it won't improve our Machine learning model , D is the closest answer , also it didn't say export it said Save , which could be manually moving the data to BQ

upvoted 2 times

victory108 3 years, 1 month ago

D. Save a history of recommendations and results of the recommendations in BigQuery, to be used as training data.

upvoted 2 times

LoveT 3 years, 1 month ago

"training data" is the key in option "D" and that's the answer

upvoted 4 times

vincy2202 3 years, 1 month ago

D seems to be the correct answer

upvoted 1 times

Question #93 *Topic 1*

A development team at your company has created a dockerized HTTPS web application. You need to deploy the application on Google Kubernetes Engine (GKE) and make sure that the application scales automatically.
How should you deploy to GKE?

A. Use the Horizontal Pod Autoscaler and enable cluster autoscaling. Use an Ingress resource to load-balance the HTTPS traffic.

B. Use the Horizontal Pod Autoscaler and enable cluster autoscaling on the Kubernetes cluster. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.

C. Enable autoscaling on the Compute Engine instance group. Use an Ingress resource to load-balance the HTTPS traffic.

D. Enable autoscaling on the Compute Engine instance group. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.

Suggested Answer: *B*
Reference: https://cloud.google.com/kubernetes-engine/docs/how-to/cluster-autoscaler

Community vote distribution
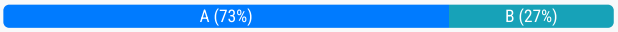

crypt0 6 years, 10 months ago
Why not using Ingress? (A)
upvoted 30 times

crypt0 6 years, 10 months ago
Why not using Ingress? (A)I think A is OK:It is A, K8s best way to LB is Ingress."Ingress is a Kubernetes resource that encapsulates a collection of rules and configuration for routing external HTTP(S) traffic to internal services. On GKE, Ingress is implemented using Cloud Load Balancing. When you create an Ingress in your cluster, GKE creates an HTTP(S) load balancer and configures it to route traffic to your application." Are you exposing multiple services through single IP address? Hence, do you need routing your traffic? Correct answer is B.My bad, as stated by other, Service doesn't support L7 load balancing. Hence, need to setup ingress resource. Correct answer is A.B is ok. https://cloud.google.com/kubernetes-engine/docs/tutorials/hello-appservice resource does a NLB using IP address, however, Ingress does HTTP(S) Load balancer. A should be an answer.
upvoted 30 times

techalik 5 years, 8 months ago
I think A is OK:
upvoted 2 times

nitinz 5 years, 5 months ago
It is A, K8s best way to LB is Ingress.
upvoted 5 times

Smart 6 years, 5 months ago
"Ingress is a Kubernetes resource that encapsulates a collection of rules and configuration for routing external HTTP(S) traffic to internal services.

On GKE, Ingress is implemented using Cloud Load Balancing. When you create an Ingress in your cluster, GKE creates an HTTP(S) load balancer and configures it to route traffic to your application."

Are you exposing multiple services through single IP address? Hence, do you need routing your traffic?

Correct answer is B.
upvoted 42 times

Smart 6 years, 5 months ago
"Ingress is a Kubernetes resource that encapsulates a collection of rules and configuration for routing external HTTP(S) traffic to internal services. On GKE, Ingress is implemented using Cloud Load Balancing. When you create an Ingress in your cluster, GKE creates an HTTP(S) load balancer and configures it to route traffic to your application." Are you exposing multiple services through single IP address? Hence, do you need routing your traffic? Correct answer is B.My bad, as stated by other, Service doesn't support L7 load balancing. Hence, need to setup ingress resource. Correct answer is A.B is ok. https://cloud.google.com/kubernetes-engine/docs/tutorials/hello-appservice resource does a NLB using IP address, however, Ingress does HTTP(S) Load balancer. A should be an answer.
upvoted 42 times

Smart Highly Voted 6 years, 5 months ago

My bad, as stated by other, Service doesn't support L7 load balancing. Hence, need to setup ingress resource. Correct answer is A.B is ok. https://cloud.google.com/kubernetes-engine/docs/tutorials/hello-appservice resource does a NLB using IP address, however, Ingress does HTTP(S) Load balancer. A should be an answer.

upvoted 47 times

Smart Highly Voted 6 years, 5 months ago

My bad, as stated by other, Service doesn't support L7 load balancing. Hence, need to setup ingress resource. Correct answer is A.

upvoted 47 times

tartar 6 years ago

B is ok. https://cloud.google.com/kubernetes-engine/docs/tutorials/hello-app

upvoted 9 times

tartar 6 years ago

B is ok. https://cloud.google.com/kubernetes-engine/docs/tutorials/hello-appservice resource does a NLB using IP address, however, Ingress does HTTP(S) Load balancer. A should be an answer.

upvoted 9 times

GopiSivanathan 5 years, 10 months ago

service resource does a NLB using IP address, however, Ingress does HTTP(S) Load balancer. A should be an answer.

upvoted 8 times

jcmoranp 6 years, 9 months ago

Name is service resource, it's B: https://cloud.google.com/kubernetes-engine/docs/concepts/service?hl=es-419

upvoted 14 times

Pime13 Most Recent 6 months ago

L7 vs. L4 Load Balancing: An Ingress resource is used for Layer 7 (HTTP/S) load balancing. It provides native support for SSL/TLS termination, URL-based routing, and multiple certificates on a single IP address. A Service of type LoadBalancer (Option B) typically creates a Layer 4 Network Load Balancer, which is less suited for advanced HTTPS features like path-based routing or managed SSL certificates. Holistic Scaling: To handle varying demand, you need two layers of scaling: Horizontal Pod Autoscaler (HPA): Dynamically increases or decreases the number of Pod replicas based on metrics like CPU utilization.Cluster Autoscaler: Automatically adds or removes worker nodes (VMs) in the cluster when Pods cannot be scheduled due to insufficient resources or when nodes are underutilized. Native Management: In GKE, you should manage scaling through Kubernetes-native objects (HPA/Cluster Autoscaler) rather than directly interacting with the underlying Compute Engine instance groups (Options C and D), as GKE automates the synchronization between these layers.

upvoted 1 times

Pime13 6 months ago

Selected Answer: A

L7 vs. L4 Load Balancing: An Ingress resource is used for Layer 7 (HTTP/S) load balancing. It provides native support for SSL/TLS termination, URL-based routing, and multiple certificates on a single IP address. A Service of type LoadBalancer (Option B) typically creates a Layer 4 Network Load Balancer, which is less suited for advanced HTTPS features like path-based routing or managed SSL certificates.

Holistic Scaling: To handle varying demand, you need two layers of scaling:

Horizontal Pod Autoscaler (HPA): Dynamically increases or decreases the number of Pod replicas based on metrics like CPU utilization.

upvoted 1 times

Pime13 6 months ago

Cluster Autoscaler: Automatically adds or removes worker nodes (VMs) in the cluster when Pods cannot be scheduled due to insufficient resources or when nodes are underutilized. Native Management: In GKE, you should manage scaling through Kubernetes-native objects (HPA/Cluster Autoscaler) rather than directly interacting with the underlying Compute Engine instance groups (Options C and D), as GKE automates the synchronization between these layers.

upvoted 1 times

drpandamd 7 months ago

A. It explains it right in here in the docs. https://docs.cloud.google.com/kubernetes-engine/docs/concepts/ingress

upvoted 1 times

JoseCloudEng1994 11 months, 3 weeks ago

Ingress Resource = HTTPS (layer 7) load balancer LoadBalancer Resource = Network (layer 4) load balancer. C, D = You should NEVER mess with the VMs manually if you are using GKE. Manage everything through GKE itself

upvoted 1 times

liorkale 1 year, 5 months ago

A is the correct answer. Let me explain why: For HTTPS traffic handling: Answer A uses an Ingress resource which is the better choice because: Ingress operates at L7 (application layer) Can handle SSL/TLS termination Can manage certificates More cost-effective (shares load balancer across multiple services) Provides URL-based routing Answer B uses a Service type LoadBalancer which is less optimal because: Creates a separate L4 load balancer for each service (more expensive) Cannot terminate SSL/TLS Cannot do path-based routing Just forwards TCP traffic on port 443 Both answers correctly mention: Horizontal Pod Autoscaler (for pod scaling) Cluster autoscaling (for node scaling) But the key differentiator is Ingress vs LoadBalancer, where Ingress is the better choice for HTTPS web applications due to its L7 capabilities and cost-effectiveness.

upvoted 3 times

igrek 1 year, 6 months ago

"team has created a dockerized HTTPS web application" - My understanding is that the HTTPS certificate is baked in the image and we can't do TLS termination. The app no longer operates on L7, but on L4 at port 443. Thus Ingress can't be used (because it can't understand encrypted traffic) and LoadBalancer service is the way to go.

upvoted 1 times

ryaryarya 1 year, 7 months ago

It states "A development team at your company has created a dockerized HTTPS web application". The web application already supports HTTPS based on that statement, so we don't need HTTPS termination and can use a LoadBalancer service. Who writes these questions, anyway?

upvoted 1 times

nareshthumma 1 year, 9 months ago

Answer A

upvoted 1 times

mstaicu 2 years, 1 month ago

A and B both create under the hood a Service of type LoadBalancer with external IP address. However, when it comes to http(s) traffic an ingress is the way to go because of ssl termination and for the routing options.

upvoted 2 times

huuthanhdlv 2 years, 2 months ago

C & D is clearly incorrect. B is incorrect because of this: "service of type LoadBalancer to load-balance the HTTPS traffic." GKE Service Load Balancer is L4 Network or Internal Load Balancer, does not support HTTPS traffic. Thus only A is correct.

upvoted 4 times

hitmax87 2 years, 3 months ago

The clue is HTTPS traffic. You need L7 stack. It can be achieved only through ingress controller.

upvoted 3 times

nanasenishino 2 years, 3 months ago

B A. Ingress resource: While Ingress can be used for external load balancing, it often requires additional configuration for HTTPS termination (offloading SSL from your application containers). Additionally, LoadBalancer services typically offer a simpler setup for basic external load balancing without HTTPS termination concerns. C & D. Compute Engine Instance Group Autoscaling: GKE manages its own nodes separate from Compute Engine instances. Autoscaling on a Compute Engine instance group wouldn't manage the Kubernetes pods or nodes effectively in this scenario.

upvoted 1 times

gun123 2 years, 7 months ago

B is correct

upvoted 1 times

bandegg 2 years, 7 months ago

I'm assuming B is the suggested answer because a the question doesn't state that the application should be available externally. Services allow exposing resources internally and to load balancers. However, it should be A, as the assumption would be a an external web application. https://cloud.google.com/kubernetes-engine/docs/concepts/service

upvoted 2 times

MahAli 2 years, 8 months ago

Most if the labs in Google boost skills discuss how to expose the deployment using a load balancer.

upvoted 2 times

AwsSuperTrooper 2 years, 8 months ago

https://cloud.google.com/kubernetes-engine/docs/concepts/ingress "This page provides a general overview of what Ingress for external Application Load Balancers is and how it works. Google Kubernetes Engine (GKE) provides a built-in and managed Ingress controller called GKE Ingress. This controller implements Ingress resources as Google Cloud load balancers for HTTP(S) workloads in GKE."

upvoted 2 times

Question #94 *Topic 1*

You need to design a solution for global load balancing based on the URL path being requested. You need to ensure operations reliability and end-to-end in- transit encryption based on Google best practices.
What should you do?

A. Create a cross-region load balancer with URL Maps.

B. Create an HTTPS load balancer with URL Maps.

C. Create appropriate instance groups and instances. Configure SSL proxy load balancing.

D. Create a global forwarding rule. Configure SSL proxy load balancing.

Suggested Answer: *B*
Reference: https://cloud.google.com/load-balancing/docs/https/url-map

Community vote distribution
B (100%)

victory108 Highly Voted 3 years, 9 months ago
B. Create an HTTPS load balancer with URL maps.
upvoted 14 times

betiy Highly Voted 5 years, 1 month ago
URL paths supported only in HTTP(S) Load balancing https://cloud.google.com/load-balancing/docs/ssl/#FAQ
upvoted 10 times

examch Most Recent 2 years, 1 month ago
B is the correct Answer, Google Cloud HTTP(S) load balancers and Traffic Director use a Google Cloud configuration resource called a URL map to route HTTP(S) requests to backend services or backend buckets. For example, with an external HTTP(S) load balancer, you can use a single URL map to route requests to different destinations based on the rules configured in the URL map: Requests for https://example.com/video go to one backend service. Requests for https://example.com/audio go to a different backend service. Requests for https://example.com/images go to a Cloud Storage backend bucket. Requests for any other host and path combination go to a default backend service. URL maps are used with the following Google Cloud products: External HTTP(S) Load Balancing (global, regional, and classic modes) Internal HTTP(S) Load Balancing Traffic Director https://cloud.google.com/load-balancing/docs/url-map-concepts
upvoted 8 times

omermahgoub 2 years, 1 month ago
B. Create an HTTPS load balancer with URL Maps. An HTTPS load balancer is a type of load balancer that can distribute incoming HTTPS traffic to one or more back-end services, such as Compute Engine instances or Google Kubernetes Engine clusters. It can also provide SSL/TLS termination, enabling you to use your own SSL/TLS certificates and keys. You can use URL Maps to configure the HTTPS load balancer to route traffic based on the URL path being requested. This allows you to set up different URL paths to be served by different back-end services, providing a high level of flexibility in your load balancing configuration.Option A, creating a cross-region load balancer with URL Maps, is also a valid solution, but it is not specifically designed for end-to-end in-transit encryption. Option C, creating appropriate instance groups and instances and configuring SSL proxy load balancing, is not a complete solution for global load balancing. SSL proxy load balancing is a feature that enables you to terminate SSL/TLS connections at the load balancer and establish a new SSL/TLS connection between the load balancer and the back-end service. It is not a global load balancing solution in and of itself. Option D, creating a global forwarding rule and configuring SSL proxy load balancing, is not a complete solution for global load balancing based on the URL path being requested. A global forwarding rule is a type of load balancing configuration that directs traffic to a specific back-end service based on the IP address and port of the incoming request. It does not allow for routing based on the URL path.
Regenerate
upvoted 1 times

omermahgoub 2 years, 1 month ago
B. Create an HTTPS load balancer with URL Maps.

An HTTPS load balancer is a type of load balancer that can distribute incoming HTTPS traffic to one or more back-end services, such as Compute Engine instances or Google Kubernetes Engine clusters. It can also provide SSL/TLS termination, enabling you to use your own SSL/TLS certificates and keys.

You can use URL Maps to configure the HTTPS load balancer to route traffic based on the URL path being requested. This allows you to set up different URL paths to be served by different back-end services, providing a high level of flexibility in your load balancing configuration.

upvoted 1 times

omermahgoub 2 years, 1 month ago

Option A, creating a cross-region load balancer with URL Maps, is also a valid solution, but it is not specifically designed for end-to-end in-transit encryption. Option C, creating appropriate instance groups and instances and configuring SSL proxy load balancing, is not a complete solution for global load balancing. SSL proxy load balancing is a feature that enables you to terminate SSL/TLS connections at the load balancer and establish a new SSL/TLS connection between the load balancer and the back-end service. It is not a global load balancing solution in and of itself. Option D, creating a global forwarding rule and configuring SSL proxy load balancing, is not a complete solution for global load balancing based on the URL path being requested. A global forwarding rule is a type of load balancing configuration that directs traffic to a specific back-end service based on the IP address and port of the incoming request. It does not allow for routing based on the URL path. Regenerate

upvoted 3 times

TonytheTiger 2 years, 2 months ago

Answer B: URL maps used with global external HTTP(S) load balancers and regional external HTTP(S) load balancer support several advanced traffic management features such as header-based traffic steering, weight-based traffic splitting, and request mirroring. https://cloud.google.com/load-balancing/docs/https#url-maps

upvoted 3 times

megumin 2 years, 3 months ago

B is ok

upvoted 1 times

Sbgani 2 years, 5 months ago

UrlMaps are used to route requests to a backend service based on rules that you define for the host and path of an incoming URL.

upvoted 5 times

Sbgani 2 years, 5 months ago

https://registry.terraform.io/providers/hashicorp/google/latest/docs/resources/compute_url_map ANS B

upvoted 2 times

DrishaS4 2 years, 6 months ago

https://cloud.google.com/load-balancing/docs/https/url-mapthank you for pointing the link

upvoted 4 times

DrishaS4 2 years, 6 months ago

Selected Answer: B

https://cloud.google.com/load-balancing/docs/https/url-map

upvoted 4 times

AzureDP900 2 years, 4 months ago

thank you for pointing the link

upvoted 1 times

haroldbenites 3 years, 2 months ago

Go for B

upvoted 2 times

vincy2202 3 years, 2 months ago

B is correct answer

upvoted 2 times

nqthien041292 3 years, 2 months ago

Vote B

upvoted 1 times

un 3 years, 9 months ago

B is correct

upvoted 1 times

ccmcwolf 3 years, 10 months ago

there are interl https load balancers they are regional https://cloud.google.com/load-balancing/docs/l7-internal

upvoted 1 times

Ausias18 3 years, 10 months ago

Answer is B

upvoted 2 times

bnlcnd 4 years ago

confused with A vs B. A has the word "cross region" but finally find out HTTP/S Load Balancing is naturally global. - B

upvoted 5 times

doumx 4 years, 2 months ago

B easy

upvoted 2 times

Question #95 *Topic 1*

You have an application that makes HTTP requests to Cloud Storage. Occasionally the requests fail with HTTP status codes of 5xx and 429. How should you handle these types of errors?

A. Use gRPC instead of HTTP for better performance.

B. Implement retry logic using a truncated exponential backoff strategy.

C. Make sure the Cloud Storage bucket is multi-regional for geo-redundancy.

D. Monitor https://status.cloud.google.com/feed.atom and only make requests if Cloud Storage is not reporting an incident.

Suggested Answer: *B*

Reference: https://cloud.google.com/storage/docs/json_api/v1/status-codes

Community vote distribution

B (100%)

bigob4ek Highly Voted 5 years, 3 months ago

Answer is B You should use exponential backoff to retry your requests when receiving errors with 5xx or 429 response codes from Cloud Storage. https://cloud.google.com/storage/docs/request-rateI agree with you, B should be rightIt is B

upvoted 44 times

bigob4ek Highly Voted 5 years, 3 months ago

Answer is B You should use exponential backoff to retry your requests when receiving errors with 5xx or 429 response codes from Cloud Storage. https://cloud.google.com/storage/docs/request-rate

upvoted 44 times

AzureDP900 2 years, 4 months ago

I agree with you, B should be right

upvoted 1 times

nitinz 3 years, 11 months ago

It is B

upvoted 1 times

Sbgani 2 years, 5 months ago

HTTP 408, 429, and 5xx response codes. Exponential backoff algorithm For requests that meet both the response and idempotency criteria, you should generally use truncated exponential backoff. Truncated exponential backoff is a standard error handling strategy for network applications in which a client periodically retries a failed request with increasing delays between requests. An exponential backoff algorithm retries requests exponentially, increasing the waiting time between retries up to a maximum backoff time. See the following workflow example to learn how exponential backoff works: You make a request to Cloud Storage. If the request fails, wait 1 + random_number_milliseconds seconds and retry the request. If the request fails, wait 2 + random_number_milliseconds seconds and retry the request. If the request fails, wait 4 + random_number_milliseconds seconds and retry the request. And so on, up to a maximum_backoff time. Continue waiting and retrying up to a maximum amount of time (deadline), but do not increase the maximum_backoff wait period between retries

upvoted 13 times

omermahgoub Most Recent 2 years, 1 month ago

. Implement retry logic using a truncated exponential backoff strategy.

HTTP status codes of 5xx and 429 typically indicate that there is a temporary issue with the service or that the rate of requests is too high. To handle these types of errors, it is generally recommended to implement retry logic in your application using a truncated exponential backoff strategy.

Truncated exponential backoff involves retrying the request after an initial delay, and then increasing the delay exponentially for each subsequent retry up to a maximum delay. This approach helps to reduce the number of failed requests and can improve the reliability of your application.

upvoted 4 times

omermahgoub 2 years, 1 month ago

. Implement retry logic using a truncated exponential backoff strategy. HTTP status codes of 5xx and 429 typically indicate that there is a temporary issue with the service or that the rate of requests is too high. To handle these types of errors, it is generally recommended to implement retry logic in your application using a truncated exponential backoff strategy. Truncated exponential backoff involves retrying the request after an initial delay, and

then increasing the delay exponentially for each subsequent retry up to a maximum delay. This approach helps to reduce the number of failed requests and can improve the reliability of your application.Option A, using gRPC instead of HTTP for better performance, is not directly related to handling HTTP status codes of 5xx and 429. gRPC is a high-performance RPC framework that can be used in place of HTTP, but it is not a solution for handling errors. Option C, making sure the Cloud Storage bucket is multi-regional for geo-redundancy, may help improve the reliability of the service, but it is not a solution for handling errors. Option D, monitoring https://status.cloud.google.com/feed.atom and only making requests if Cloud Storage is not reporting an incident, is not a practical solution for handling errors. This approach would require constantly monitoring the status page and could result in significant delays in processing requests. Instead, it is generally recommended to implement retry logic in your application to handle errors.

upvoted 4 times

omermahgoub 2 years, 1 month ago

Option A, using gRPC instead of HTTP for better performance, is not directly related to handling HTTP status codes of 5xx and 429. gRPC is a high-performance RPC framework that can be used in place of HTTP, but it is not a solution for handling errors. Option C, making sure the Cloud Storage bucket is multi-regional for geo-redundancy, may help improve the reliability of the service, but it is not a solution for handling errors. Option D, monitoring https://status.cloud.google.com/feed.atom and only making requests if Cloud Storage is not reporting an incident, is not a practical solution for handling errors. This approach would require constantly monitoring the status page and could result in significant delays in processing requests. Instead, it is generally recommended to implement retry logic in your application to handle errors.

upvoted 2 times

megumin 2 years, 3 months ago

B is ok

upvoted 1 times

Sbgani 2 years, 5 months ago

https://cloud.google.com/storage/docs/retry-strategy

upvoted 2 times

DrishaS4 2 years, 6 months ago

2xx – successful requests; 4xx, 5xx – failed requests; 3xx – requests that require redirect.
https://cloud.google.com/storage/docs/json_api/v1/status-codes

upvoted 2 times

haroldbenites 3 years, 2 months ago

Go for B

upvoted 2 times

vincy2202 3 years, 2 months ago

B is the correct answer

upvoted 2 times

nqthien041292 3 years, 2 months ago

Vote B

upvoted 2 times

joe2211 3 years, 2 months ago

vote B

upvoted 1 times

MaxNRG 3 years, 3 months ago

B – Implement retry logic using a truncated exponential backoff strategy. Per HTTP status and error codes for JSON the status codes are: 2xx – successful requests; 4xx, 5xx – failed requests; 3xx – requests that require redirect. https://cloud.google.com/storage/docs/json_api/v1/status-codes 429 – Too many requests: your app tries to use more that its limit, additional requests will fail. Decrease your client's requests and/or use truncated exponential backoff (used for all requests with 5xx and 429 errors). https://cloud.google.com/storage/docs/retry-strategy

upvoted 2 times

victory108 3 years, 9 months ago

B. Use Deployment Manager to automate service provisioning. Use Stackdriver to monitor and debug your tests.

upvoted 2 times

victory108 3 years, 9 months ago

B. Use Deployment Manager to automate service provisioning. Use Stackdriver to monitor and debug your tests.This B. Implement retry logic using a truncated exponential backoff strategy.

upvoted 2 times

victory108 3 years, 9 months ago

This B. Implement retry logic using a truncated exponential backoff strategy.

upvoted 1 times

un 3 years, 9 months ago

Answer is B. Link provided by bigob4ek has details

upvoted 1 times

Ausias18 3 years, 10 months ago

Answer is B

upvoted 1 times

CloudGenious 4 years ago

As per google, if you run into any issue as increase latency or erroe rate ,pause your ramp up this give cloudstorage more time to scale your bucket .

Best is backoff when 5xx ,429,408 response code

upvoted 2 times

bnlcnd 4 years ago

https://cloud.google.com/storage/docs/exponential-backoff - B

upvoted 2 times

awadheshk 4 years, 4 months ago

B is correct

upvoted 1 times

Question #96 *Topic 1*

You need to develop procedures to test a disaster plan for a mission-critical application. You want to use Google-recommended practices and native capabilities within GCP.
What should you do?

A. Use Deployment Manager to automate service provisioning. Use Activity Logs to monitor and debug your tests.

B. Use Deployment Manager to automate service provisioning. Use Stackdriver to monitor and debug your tests.

C. Use gcloud scripts to automate service provisioning. Use Activity Logs to monitor and debug your tests.

D. Use gcloud scripts to automate service provisioning. Use Stackdriver to monitor and debug your tests.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (88%) D (12%)

crypt0 Highly Voted 5 years, 9 months ago
I think answer B is correct: https://cloud.google.com/solutions/dr-scenarios-planning-guideB is okB is correcthttps://cloud.google.com/solutions/dr-scenarios-planning-guide#test_your_plan_regularlyIt is B, Google Best practice ---> never use scripts. They do not trust anyone else's code it seems.
upvoted 54 times

crypt0 Highly Voted 5 years, 9 months ago
I think answer B is correct: https://cloud.google.com/solutions/dr-scenarios-planning-guide
upvoted 54 times

tartar 5 years ago
B is ok
upvoted 11 times

kumarp6 4 years, 9 months ago
B is correct
upvoted 2 times

fraloca 4 years, 7 months ago
https://cloud.google.com/solutions/dr-scenarios-planning-guide#test_your_plan_regularly
upvoted 2 times

nitinz 4 years, 5 months ago
It is B, Google Best practice ---> never use scripts. They do not trust anyone else's code it seems.
upvoted 12 times

passnow 5 years, 8 months ago
Boom, everyone studied and did their labs, stackdriver is google's recommended tool for monitoring and debbuging. I agree with u all that B is the correct answer
upvoted 23 times

duzapo Most Recent 1 year, 11 months ago
D its correct cause are 3 multiregions-availables and one bucket only can deploy in one multi region
https://cloud.google.com/storage/docs/locations?hl=es-419#location-mr
upvoted 1 times

jalberto 2 years ago
I think B is the correct answer
upvoted 1 times

faridomu 2 years, 2 months ago
Why not A?
upvoted 2 times

Jlharidon 2 years, 8 months ago
Deploy managment + Stackdriver trained ig GCSB
upvoted 1 times

megumin 2 years, 9 months ago

B is ok

upvoted 1 times

AzureDP900 2 years, 10 months ago

B is right

upvoted 1 times

tycho 3 years ago

in practice, D could work as well..

upvoted 1 times

tycho 3 years ago

in practice, D could work as well..yeah, but only native solutions should be taken into consideration (as stated in requirements), so scripts are basically ruled out

upvoted 1 times

adacek1 2 years, 6 months ago

yeah, but only native solutions should be taken into consideration (as stated in requirements), so scripts are basically ruled out

upvoted 1 times

gaojun 3 years, 4 months ago

Answer B is correct

upvoted 1 times

Skr6266 3 years, 5 months ago

Deployment Manager + Cloud Monitoring and Logging solution.

upvoted 1 times

haroldbenites 3 years, 8 months ago

Go for B

upvoted 1 times

vincy2202 3 years, 8 months ago

B is the correct answer

upvoted 1 times

nqthien041292 3 years, 8 months ago

Vote B

upvoted 1 times

ganeshrev 3 years, 9 months ago

Google recommended Practice

upvoted 1 times

victory108 4 years, 3 months ago

B. Use Deployment Manager to automate service provisioning. Use Stackdriver to monitor and debug your tests.

upvoted 4 times

un 4 years, 3 months ago

B is correct

upvoted 1 times

Question #97 *Topic 1*

Your company creates rendering software which users can download from the company website. Your company has customers all over the world. You want to minimize latency for all your customers. You want to follow Google-recommended practices.
How should you store the files?

A. Save the files in a Multi-Regional Cloud Storage bucket.

B. Save the files in a Regional Cloud Storage bucket, one bucket per zone of the region.

C. Save the files in multiple Regional Cloud Storage buckets, one bucket per zone per region.

D. Save the files in multiple Multi-Regional Cloud Storage buckets, one bucket per multi-region.

Suggested Answer: *A*
Official Answer: A

Community vote distribution



JoeShmoe Highly Voted 6 years, 9 months ago
Its D, create multi region buckets in Americas, Europe and Asia
upvoted 65 times

JoeShmoe Highly Voted 6 years, 9 months ago
Its D, create multi region buckets in Americas, Europe and AsiaThis can't be D. It should be A.wrong its Awhy " multiple Multi-Regional"? - A should be the right ans & addressing the global users - "More importantly, is that multiregional heavily leverages Edge caching and CDNs to provide the content to the end user" https://medium.com/google-cloud/google-cloud-storage-what-bucket-class-for-the-best-performance-5c847ac8f9f2because a multi-regional includes all the locations of ONE region, not the others.What is point of Multi-Regional bucket, if this need to saved multiple times. I believe option (D) is for creating confusion only. It should be (A)..Check the current create bucket UI. You cannot select Asia multi-region and US multi-region at the same go. So to support global customer, you need to create multiple Multi-region buckets.Read the question again.. I think (d) is correct.. eg. 1 bucket in US-multi-region, 2nd in AS-multi-region, 3rd in EU-multi-regionYes, D seems correct. There are 3 multi-regions: ASIA, EU and US. In order to be global, there must be multi-region buckets in this 3 locations. Reference: https://cloud.google.com/storage/docs/locations#location-mrLet's try option A: you select a single multi-region bucket (e.g. Americas). Are you improving the latency of your clients in Asia? You do not. Thus, Option A is not complete.
upvoted 65 times

Urban_Life 4 years, 8 months ago
This can't be D. It should be A.
upvoted 8 times

kilo10x 3 years, 1 month ago
wrong its A
upvoted 2 times

MyPractice 6 years, 7 months ago
why " multiple Multi-Regional"? - A should be the right ans & addressing the global users - "More importantly, is that multiregional heavily leverages Edge caching and CDNs to provide the content to the end user" https://medium.com/google-cloud/google-cloud-storage-what-bucket-class-for-the-best-performance-5c847ac8f9f2because a multi-regional includes all the locations of ONE region, not the others.
upvoted 14 times

MyPractice 6 years, 7 months ago
why " multiple Multi-Regional"? - A should be the right ans & addressing the global users - "More importantly, is that multiregional heavily leverages Edge caching and CDNs to provide the content to the end user" https://medium.com/google-cloud/google-cloud-storage-what-bucket-class-for-the-best-performance-5c847ac8f9f2
upvoted 14 times

xavi1 4 years, 12 months ago
because a multi-regional includes all the locations of ONE region, not the others.
upvoted 13 times

AmitAr 4 years, 3 months ago
What is point of Multi-Regional bucket, if this need to saved multiple times. I believe option (D) is for creating confusion only. It should be (A)..

upvoted 19 times

AmitAr 4 years, 3 months ago

What is point of Multi-Regional bucket, if this need to saved multiple times. I believe option (D) is for creating confusion only. It should be (A)..Check the current create bucket UI. You cannot select Asia multi-region and US multi-region at the same go. So to support global customer, you need to create multiple Multi-region buckets.Read the question again.. I think (d) is correct.. eg. 1 bucket in US-multi-region, 2nd in AS-multi-region, 3rd in EU-multi-regionYes, D seems correct. There are 3 multi-regions: ASIA, EU and US. In order to be global, there must be multi-region buckets in this 3 locations. Reference: https://cloud.google.com/storage/docs/locations#location-mrLet's try option A: you select a single multi-region bucket (e.g. Americas). Are you improving the latency of your clients in Asia? You do not. Thus, Option A is not complete.

upvoted 19 times

turbo8p 3 years, 9 months ago

Check the current create bucket UI. You cannot select Asia multi-region and US multi-region at the same go. So to support global customer, you need to create multiple Multi-region buckets.

upvoted 13 times

AmitAr 4 years, 3 months ago

Read the question again.. I think (d) is correct.. eg. 1 bucket in US-multi-region, 2nd in AS-multi-region, 3rd in EU-multi-regionYes, D seems correct. There are 3 multi-regions: ASIA, EU and US. In order to be global, there must be multi-region buckets in this 3 locations. Reference: https://cloud.google.com/storage/docs/locations#location-mr

upvoted 10 times

AmitAr 4 years, 3 months ago

Read the question again.. I think (d) is correct.. eg. 1 bucket in US-multi-region, 2nd in AS-multi-region, 3rd in EU-multi-region

upvoted 10 times

giovanicascaes 3 years, 4 months ago

Yes, D seems correct. There are 3 multi-regions: ASIA, EU and US. In order to be global, there must be multi-region buckets in this 3 locations. Reference: https://cloud.google.com/storage/docs/locations#location-mr

upvoted 7 times

JaimeMS 2 years, 2 months ago

Let's try option A: you select a single multi-region bucket (e.g. Americas). Are you improving the latency of your clients in Asia? You do not. Thus, Option A is not complete.

upvoted 2 times

MeasService 6 years, 10 months ago

I would go with A (https://cloud.google.com/storage/docs/locations)

upvoted 36 times

eff12c1 Most Recent 3 months ago

A multi-region Cloud Storage bucket stores data across multiple regions within a broad geography, such as US, EU, or ASIA, and is designed for higher availability and broad user access compared with a single regional bucket. Google describes multi-region as a large geographic area containing multiple regions.

upvoted 2 times

CyberPeter 6 months, 2 weeks ago

A matches Google-recommended practices Google's own guidance for globally accessed static content: Use Cloud Storage multi-regional buckets Optionally front with Cloud CDN (not mentioned in the question) This gives: Low latency worldwide High availability Simple architecture

upvoted 1 times

kurahaz 8 months, 1 week ago

Why A is Correct Purpose-Built Class: The Multi-Regional storage class is specifically designed for storing data that is frequently accessed ("hot" data) such as serving website content, video, or mobile application data to users distributed across large geographic areas. Geo-Redundancy: It stores your data redundantly across at least two geographically separated regions (e.g., within the US, EU, or Asia). This ensures the highest availability.

upvoted 2 times

KCjoe 8 months, 3 weeks ago

You don't need mutiple "Multi-Regional bucket" , you only need one.

upvoted 2 times

lokiinaction 1 year, 2 months ago

When you want to minimize latency for accessing content within Cloud Storage, the best Google Cloud practice is to choose Regional Storage, but strategically, in the region closest to your primary users or compute resources

upvoted 1 times

MrAZ105 1 year, 3 months ago

A is the best answer because a Multi-Regional Cloud Storage bucket offers low latency, high availability, and global access, aligning with Google's best practices for user-facing content.

upvoted 3 times

francisco94 1 year, 4 months ago

Why not the other options? B: Regional buckets only store data in a single region — this would increase latency for users far from that region. C: You can't create buckets per zone; Cloud Storage is regional or multi-regional, not zonal. D: Managing multiple multi-regional buckets is unnecessary and adds complexity without performance benefits.

upvoted 2 times

apb98 1 year, 4 months ago

It's option D. Is not Google best practice, but it's the only way to make it global as the answers don't mention the use of Multi-Regional bucket + Cloud CDN or edge caching. So D is the only one.

upvoted 1 times

OnoPa 1 year, 4 months ago

While Option D would work, it is not needed. A single multi-regional bucket handles this case, and is the google recomended practice. Also, it is more complex to manage multiple buckets.

upvoted 3 times

halifax 1 year, 5 months ago

When you need to serve content to users globally with minimal latency, a Multi-Regional Cloud Storage bucket is the recommended solution. It handles the distribution and redundancy for you.

upvoted 3 times

kahinah 1 year, 5 months ago

Multi-regional buckets replicate data across geographically dispersed regions within a continent (e.g., ASIA, EU, or US multi-regions), ensuring users download from the nearest available location

upvoted 3 times

mselmi 1 year, 5 months ago

A multi-region is a large geographic area, such as the United States, that contains two or more geographic places.
https://cloud.google.com/storage/docs/locations

upvoted 1 times

Amrx 1 year, 9 months ago

D is correct. Multi-region buckets are still specific to their own regional area, Americas, Europe and Asia. It's not A, doesn't cover the whole world.

upvoted 1 times

hehe_24 1 year, 10 months ago

I go for A. I never came across"multiple multi-region"

upvoted 1 times

JaimeMS 2 years, 2 months ago

Its D, create multi region buckets in Americas, Europe and Asia

upvoted 2 times

Question #98 *Topic 1*

Your company acquired a healthcare startup and must retain its customers' medical information for up to 4 more years, depending on when it was created. Your corporate policy is to securely retain this data, and then delete it as soon as regulations allow.
Which approach should you take?

A. Store the data in Google Drive and manually delete records as they expire.

B. Anonymize the data using the Cloud Data Loss Prevention API and store it indefinitely.

C. Store the data in Cloud Storage and use lifecycle management to delete files when they expire.

D. Store the data in Cloud Storage and run a nightly batch script that deletes all expired data.

Suggested Answer: *C*
Official Answer: C

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

AWS56 Highly Voted 6 years, 7 months ago
Agree C
upvoted 24 times

JoseCloudEng1994 Most Recent 11 months, 3 weeks ago
C beyond the shadow of the doubt. As soon as you hear automatic deletion = Cloud Storage. In general Google does not like or advertise to create scripts that do stuff to their systems so between C or D then C is 100% the correct answer
upvoted 1 times

Abhi_ma_cdl_22 1 year, 4 months ago
C is the right answer
upvoted 1 times

desertlotus1211 1 year, 8 months ago
Selected Answer: B
Why not B? It's patients' health records...
upvoted 1 times

desertlotus1211 1 year, 8 months ago
Why not B? It's patients' health records...Because the requirement says 4 more years retention, not indefinitely.
upvoted 1 times

Peto12 1 year, 7 months ago
Because the requirement says 4 more years retention, not indefinitely.
upvoted 1 times

i_maddog_i 3 years, 5 months ago
It's C
upvoted 1 times

megumin 3 years, 9 months ago
C is ok
upvoted 1 times

AzureDP900 3 years, 10 months ago
I agree with C
upvoted 1 times

DrishaS4 4 years ago
go for C
upvoted 1 times

Dhiraj03 4 years, 2 months ago
Options C undoubtedly
upvoted 1 times

gaojun 4 years, 4 months ago

Go for C

upvoted 1 times

Rajasa 4 years, 7 months ago

Go for C

upvoted 1 times

haroldbenites 4 years, 8 months ago

Go for C

upvoted 1 times

vincy2202 4 years, 8 months ago

C is the correct answer

upvoted 1 times

gabrielzeven 4 years, 8 months ago

D sounds like i would do it, but C sound like a lab or exam

upvoted 1 times

nqthien041292 4 years, 8 months ago

Vote C

upvoted 1 times

victory108 5 years, 3 months ago

C. Store the data in Cloud Storage and use lifecycle management to delete files when they expire.

upvoted 4 times

un 5 years, 3 months ago

C is correct

upvoted 1 times

Question #99 *Topic 1*

You are deploying a PHP App Engine Standard service with Cloud SQL as the backend. You want to minimize the number of queries to the database.
What should you do?

A. Set the memcache service level to dedicated. Create a key from the hash of the query, and return database values from memcache before issuing a query to Cloud SQL.

B. Set the memcache service level to dedicated. Create a cron task that runs every minute to populate the cache with keys containing query results.

C. Set the memcache service level to shared. Create a cron task that runs every minute to save all expected queries to a key called ɹ€cached_queriesɹ€.

D. Set the memcache service level to shared. Create a key called ɹ€cached_queriesɹ€, and return database values from the key before using a query to Cloud SQL.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (100%)

hiteshrup Highly Voted 4 years, 8 months ago
A dedicated memset is always better than shared until cost-effectiveness specify in the exam as objective. So Option C and D are ruled out. From A and B, Option B is sending and updating query every minutes which is over killing. So reasonable option left with A which balance performance and cost. My answer will be AGood job bro
upvoted 31 times

hiteshrup Highly Voted 4 years, 8 months ago
A dedicated memset is always better than shared until cost-effectiveness specify in the exam as objective. So Option C and D are ruled out.

From A and B, Option B is sending and updating query every minutes which is over killing. So reasonable option left with A which balance performance and cost.

My answer will be A
upvoted 31 times

ArtistS 1 year, 9 months ago
Good job bro
upvoted 1 times

Eroc 5 years, 9 months ago
https://cloud.google.com/appengine/docs/standard/php/memcache/using
upvoted 24 times

Eroc 5 years, 9 months ago
https://cloud.google.com/appengine/docs/standard/php/memcache/usingA is okhttps://cloud.google.com/memorystore/docs/redis/redis-overviewA is correct
upvoted 24 times

tartar 5 years ago
A is ok
upvoted 11 times

dlzhang 4 years, 2 months ago
https://cloud.google.com/memorystore/docs/redis/redis-overview
upvoted 2 times

nitinz 4 years, 5 months ago
A is correct
upvoted 6 times

francisco94 Most Recent 1 year, 4 months ago

Since Memcache is deprecated in newer runtimes, you should: Use Memorystore for Redis, a managed and scalable caching layer. There's no "service level" to set anymore like with the legacy memcache — Redis handles performance and availability via tier selection (Basic or Standard).

upvoted 3 times

Abhi_ma_cdl_22 1 year, 4 months ago

A is correct

upvoted 1 times

Sur_Nikki 2 years, 3 months ago

Best is A

upvoted 2 times

megumin 2 years, 9 months ago

A is ok

upvoted 1 times

AzureDP900 2 years, 10 months ago

A is fine.. dedicated mem cache

upvoted 1 times

DrishaS4 3 years ago

https://cloud.google.com/appengine/docs/standard/php/memcache/using

upvoted 1 times

gaojun 3 years, 4 months ago

Obviously, the answer is A

upvoted 1 times

ehgm 3 years, 7 months ago

Dedicated and shared will resolve the problem, the key is: store all queries in only one key "cached_queries" is not good, we have limits: https://cloud.google.com/appengine/docs/standard/python/memcache Create a key of each query is better.

upvoted 4 times

vincy2202 3 years, 8 months ago

A is the correct answer

upvoted 2 times

nqthien041292 3 years, 8 months ago

Vote A

upvoted 1 times

joe2211 3 years, 8 months ago

vote A

upvoted 1 times

victory108 4 years, 3 months ago

A. Set the memcache service level to dedicated. Create a key from the hash of the query, and return database values from memcache before issuing a query to Cloud SQL.

upvoted 3 times

un 4 years, 3 months ago

A is correct

upvoted 1 times

Ausias18 4 years, 4 months ago

Answer is A

upvoted 1 times

ga 4 years, 5 months ago

A is correct

upvoted 1 times

Question #100 *Topic 1*

You need to ensure reliability for your application and operations by supporting reliable task scheduling for compute on GCP. Leveraging Google best practices, what should you do?

A. Using the Cron service provided by App Engine, publish messages directly to a message-processing utility service running on Compute Engine instances.

B. Using the Cron service provided by App Engine, publish messages to a Cloud Pub/Sub topic. Subscribe to that topic using a message-processing utility service running on Compute Engine instances.

C. Using the Cron service provided by Google Kubernetes Engine (GKE), publish messages directly to a message-processing utility service running on Compute Engine instances.

D. Using the Cron service provided by GKE, publish messages to a Cloud Pub/Sub topic. Subscribe to that topic using a message-processing utility service running on Compute Engine instances.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (97%) (3

JoeShmoe Highly Voted 6 years, 3 months ago
Answer is B
upvoted 32 times

Smart Highly Voted 5 years, 11 months ago
B is correct. More appropriately: https://cloud.google.com/solutions/reliable-task-scheduling-compute-engine
upvoted 30 times

Smart 5 years, 11 months ago
B is correct. More appropriately: https://cloud.google.com/solutions/reliable-task-scheduling-compute-enginehttps://cloud.google.com/solutions/reliable-task-scheduling-compute-engine#schedule-compute-engine
upvoted 30 times

fraloca 5 years, 1 month ago
https://cloud.google.com/solutions/reliable-task-scheduling-compute-engine#schedule-compute-engine
upvoted 4 times

xaqanik 2 years ago
You can create Cron job using HTTP endpoint, Pub/Sub and App engine.
upvoted 1 times

odacir 2 years, 3 months ago
Selected Answer: B
Answer is B, but this question is outdated, Today the best practices for cron is Cloud Scheduler: fully managed enterprise-grade cron job scheduler https://cloud.google.com/scheduler/?gad_source=1&gclsrc=ds&gclsrc=ds
upvoted 22 times

odacir 2 years, 3 months ago
Answer is B, but this question is outdated, Today the best practices for cron is Cloud Scheduler: fully managed enterprise-grade cron job scheduler https://cloud.google.com/scheduler/?gad_source=1&gclsrc=ds&gclsrc=dsThanks... I was a little confused by this options
upvoted 22 times

JaimeMS Most Recent 1 year, 8 months ago
Thanks... I was a little confused by this options
upvoted 2 times

JPA210 2 years, 4 months ago
This seems to be an old question, despite B could be the more correct answer, it is not exactly a good one. 'Using the Cron service provided by App Engine', the cron service is provided by Cloud Scheduler, not App Engine. App Engine HTTP endpoint can be a target for the cron task.
upvoted 10 times

salim_ 2 years, 9 months ago

https://cloud.google.com/blog/products/gcp/reliable-task-scheduling-on-google-compute-engine
upvoted 3 times

rr4444 2 years, 11 months ago
Something feels missing/broken about this question Even before comments in discussion that correctly mentioned Cloud Scheduler, which is not mentioned in the questionThis is because cloud scheduler is a newly released service which is a replacement to cloud app engine cron service.
upvoted 6 times

rr4444 2 years, 11 months ago
Something feels missing/broken about this question

Even before comments in discussion that correctly mentioned Cloud Scheduler, which is not mentioned in the question
upvoted 6 times

parthkulkarni998 2 years, 2 months ago
This is because cloud scheduler is a newly released service which is a replacement to cloud app engine cron service.
upvoted 1 times

dataqueen_3110 3 years, 1 month ago
"By using Cloud Scheduler for scheduling and Pub/Sub for distributed messaging, you can build an application to reliably schedule tasks across a fleet of Compute Engine instances." https://cloud.google.com/architecture/reliable-task-scheduling-compute-engine Answer is B. (Note: It was down to B or D but containerization was not mentioned)
upvoted 2 times

beehive 3 years, 1 month ago
Answer is B. Cloud Scheduler provides a fully managed, enterprise-grade service that lets you schedule events. After you have scheduled a job, Cloud Scheduler will call the configured event handlers, which can be App Engine services, HTTP endpoints, or Pub/Sub subscriptions. To run tasks on your Compute Engine instance in response to Cloud Scheduler events, you need to relay the events to those instances. One way to do this is by calling an HTTP endpoint that runs on your Compute Engine instances. Another option is to pass messages from Cloud Scheduler to your Compute Engine instances using Pub/Sub.
upvoted 8 times

habros 3 years, 2 months ago
A and C are out… messages are to be sent to pub sub and processed using a client. D is overkill for this purpose
upvoted 2 times

megumin 3 years, 3 months ago
B is ok
upvoted 1 times

AzureDP900 3 years, 4 months ago
B is right
upvoted 1 times

Nirca 3 years, 5 months ago
Ans is B https://cloud.google.com/architecture/reliable-task-scheduling-compute-engine
upvoted 1 times

Sbgani 3 years, 5 months ago
Ans is B https://cloud.google.com/architecture/reliable-task-scheduling-compute-engine refer the examples with diagram
upvoted 1 times

pp0709 3 years, 5 months ago
Selected Answer: D
This solution can be implemented using both A and D 1) With App Engine - https://cloud.google.com/appengine/docs/flexible/nodejs/scheduling-jobs-with-cron-yaml 2) With GKE - https://cloud.google.com/kubernetes-engine/docs/how-to/cronjobs

They ask for best practices and it's well known that GKE (aka containers) is the best practice for building modern infra solution.

Yet another confusing PCA question on the card. Honestly, think the quality of the questions can be mightily improved.
upvoted 1 times

pp0709 3 years, 5 months ago
This solution can be implemented using both A and D 1) With App Engine - https://cloud.google.com/appengine/docs/flexible/nodejs/scheduling-jobs-with-cron-yaml 2) With GKE - https://cloud.google.com/kubernetes-engine/docs/how-to/cronjobs They ask for best practices and it's well known that GKE (aka containers) is the best practice for building modern infra solution. Yet another confusing PCA question on the card. Honestly, think the quality of the questions can be mightily improved.A is a bad solution as "send message directly to the utility" is not really reliable, you'd want pub/sub

in between.Sorry, can be implemented using both B and DYou right, but D is overkill. So B if the best practices for this task.GKE is too expensive if all you are after is cron scheduling.

upvoted 1 times

medi01 2 years, 10 months ago

A is a bad solution as "send message directly to the utility" is not really reliable, you'd want pub/sub in between.

upvoted 1 times

pp0709 3 years, 5 months ago

Sorry, can be implemented using both B and DYou right, but D is overkill. So B if the best practices for this task.

upvoted 2 times

pp0709 3 years, 5 months ago

Sorry, can be implemented using both B and D

upvoted 2 times

kapara 2 years, 6 months ago

You right, but D is overkill. So B if the best practices for this task.

upvoted 1 times

BiddlyBdoyng 3 years, 4 months ago

GKE is too expensive if all you are after is cron scheduling.

upvoted 1 times

6721sora 3 years, 5 months ago

B says Appengine. But Cloud Scheduler is itself a managed service. To schedule jobs via AppEngine, the cron.yaml has to be used. It can be done similarly via GKE as well. This question is confusing

upvoted 2 times

DrishaS4 3 years, 6 months ago

https://cloud.google.com/solutions/reliable-task-scheduling-compute-engine

upvoted 1 times

Question #101 *Topic 1*

Your company is building a new architecture to support its data-centric business focus. You are responsible for setting up the network. Your company's mobile and web-facing applications will be deployed on-premises, and all data analysis will be conducted in GCP. The plan is to process and load 7 years of archived .csv files totaling 900 TB of data and then continue loading 10 TB of data daily. You currently have an existing 100-MB internet connection.
What actions will meet your company's needs?

A. Compress and upload both archived files and files uploaded daily using the gsutil ג€"m option.

B. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storage. Establish a connection with Google using a Dedicated Interconnect or Direct Peering connection and use it to upload files daily.

C. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storage. Establish one Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily using the gsutil ג€"m option.

D. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storage. Establish a Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (86%) C (10%) (5%)

KouShikyou Highly Voted 6 years, 10 months ago
With option A, daily data would take 27 hours. My answer is B. How do you think?
upvoted 54 times

KouShikyou Highly Voted 6 years, 10 months ago
With option A, daily data would take 27 hours. My answer is B. How do you think?Also with option A you need months to download archive files(900TB)B is correctok but dedicated connection is available from 10 GBPS right where as in question it says internet connection is 100 MB, to me D is correct.Dedicated Interconnect will be a new connection and will not run over the existing internet connection. With dedicated interconnect the existing ISP becomes irrelevant. If you were trying to use VPN the existing internet connection would be relevant. Answer is B.it is BDirect peering is meant only to connect to G Suite Services. Its reference may invalidate the whole answer.True. B is the most apt answer with just this extra bit "direct peering" raising some confusion.
upvoted 54 times

xaqanik 2 years, 6 months ago
Also with option A you need months to download archive files(900TB)
upvoted 1 times

kumarp6 5 years, 9 months ago
B is correct
upvoted 5 times

kumarp6 5 years, 9 months ago
B is correctok but dedicated connection is available from 10 GBPS right where as in question it says internet connection is 100 MB, to me D is correct.Dedicated Interconnect will be a new connection and will not run over the existing internet connection. With dedicated interconnect the existing ISP becomes irrelevant. If you were trying to use VPN the existing internet connection would be relevant. Answer is B.
upvoted 5 times

Jay_82 5 years, 8 months ago
ok but dedicated connection is available from 10 GBPS right where as in question it says internet connection is 100 MB, to me D is correct.
upvoted 5 times

Jay_82 5 years, 8 months ago
ok but dedicated connection is available from 10 GBPS right where as in question it says internet connection is 100 MB, to me D is correct.Dedicated Interconnect will be a new connection and will not run over the existing internet connection. With dedicated interconnect the existing ISP becomes irrelevant. If you were trying to use VPN the existing internet connection would be relevant. Answer is B.
upvoted 5 times

9xnine 4 years, 2 months ago

Dedicated Interconnect will be a new connection and will not run over the existing internet connection. With dedicated interconnect the existing ISP becomes irrelevant. If you were trying to use VPN the existing internet connection would be relevant. Answer is B.

upvoted 6 times

nitinz 5 years, 5 months ago

it is B

upvoted 7 times

malequardos 5 years, 3 months ago

Direct peering is meant only to connect to G Suite Services. Its reference may invalidate the whole answer.

upvoted 3 times

malequardos 5 years, 3 months ago

Direct peering is meant only to connect to G Suite Services. Its reference may invalidate the whole answer.True. B is the most apt answer with just this extra bit "direct peering" raising some confusion.

upvoted 3 times

NG123 4 years, 1 month ago

True. B is the most apt answer with just this extra bit "direct peering" raising some confusion.

upvoted 1 times

wk 6 years, 10 months ago

Agree B. 100Mbps connections for 10TB data transfer is takes too long https://cloud.google.com/solutions/transferring-big-data-sets-to-gcp#closenot 100Mbps. 100MBeven with 100MB internet it's slow. It's 800 Mbps and transfer for 10 TB will take 2 daysThere is no such thing as a "100MB" internet connection :) . That must be a speed (per second), and I would guess that the "B" is just a typo (it is highly atypical to measure bandwidth in Bps).

upvoted 21 times

wk 6 years, 10 months ago

Agree B. 100Mbps connections for 10TB data transfer is takes too long

https://cloud.google.com/solutions/transferring-big-data-sets-to-gcp#close

upvoted 21 times

JJu 6 years, 8 months ago

not 100Mbps. 100MB

upvoted 3 times

JJu 6 years, 8 months ago

not 100Mbps. 100MBeven with 100MB internet it's slow. It's 800 Mbps and transfer for 10 TB will take 2 daysThere is no such thing as a "100MB" internet connection :) . That must be a speed (per second), and I would guess that the "B" is just a typo (it is highly atypical to measure bandwidth in Bps).

upvoted 3 times

misho 6 years, 2 months ago

even with 100MB internet it's slow. It's 800 Mbps and transfer for 10 TB will take 2 days

upvoted 5 times

bogd 5 years, 5 months ago

There is no such thing as a "100MB" internet connection :) . That must be a speed (per second), and I would guess that the "B" is just a typo (it is highly atypical to measure bandwidth in Bps).

upvoted 5 times

CtrlX Most Recent 11 months, 2 weeks ago

answer D. Was in my exam

upvoted 1 times

lokiinaction 1 year, 2 months ago

direct peering can not be the option in practice, the reality is to use gsutil -m option

upvoted 1 times

halifax 1 year, 5 months ago

Option B is the correct answer, but the term "or Direct Peering connection" is probably a typo. Direct peering is GCP to GCP service NOT an on-prem to GCP service.

upvoted 1 times

plumbig11 1 year, 7 months ago

Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storage. (In this case is the more appropriated solution) Establish a connection with Google using a Dedicated Interconnect or Direct Peering connection and use it to upload files

daily. (Because of the size Dedicated interconnect or peering performs better and offer a better latency)

upvoted 2 times

RickMorais 2 years, 1 month ago

There is no option than B.

upvoted 1 times

ManojNegi 2 years, 2 months ago

B is correct

upvoted 1 times

Devx198912233 2 years, 6 months ago

B would be correct as pubsub b service might redeliver messages. When you receive messages in order and the Pub/Sub service redelivers a message with an ordering key, Pub/Sub maintains order by also redelivering the subsequent messages with the same ordering key. The Pub/Sub service redelivers these messages in the order that it originally received them.

upvoted 1 times

gun123 2 years, 7 months ago

B has an option of direct peering too which is not a recommended practise

upvoted 1 times

BisoWafik 2 years, 8 months ago

B makes sense

upvoted 1 times

odacir 2 years, 9 months ago

B. https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#time

upvoted 1 times

Arun_m_123 2 years, 10 months ago

I don't think the company needs "dedicated intereconnect" - because it is clearly said that the company wants to do "data-centric" business. Dedicated interconnect is more for having private-access to google cloud. C seems like a correct option to melt's B, transferring over 100mbps 10TB daily is not possible, because could take up to 12 days.. You need a better connection.

upvoted 1 times

Arun_m_123 2 years, 10 months ago

Selected Answer: C

I don't think the company needs "dedicated intereconnect" - because it is clearly said that the company wants to do "data-centric" business. Dedicated interconnect is more for having private-access to google cloud.

C seems like a correct option to me

upvoted 1 times

odacir 2 years, 9 months ago

It's B, transferring over 100mbps 10TB daily is not possible, because could take up to 12 days.. You need a better connection.

upvoted 1 times

telp 3 years, 5 months ago

Answer B => Dedicated interconnect will provide a private network with 10gbs. The internet limited to 100 mb is not possible to use cloud VPN (it will use public internet so be limited for the daily)

upvoted 2 times

sunny2421 3 years, 7 months ago

B is correct.

upvoted 1 times

habros 3 years, 8 months ago

B. Since it is a new network just sign up for a dedicated line…

upvoted 1 times

megumin 3 years, 9 months ago

B is ok

upvoted 1 times

Question #102 *Topic 1*

You are developing a globally scaled frontend for a legacy streaming backend data API. This API expects events in strict chronological order with no repeat data for proper processing.
Which products should you deploy to ensure guaranteed-once FIFO (first-in, first-out) delivery of data?

A. Cloud Pub/Sub alone

B. Cloud Pub/Sub to Cloud Dataflow

C. Cloud Pub/Sub to Stackdriver

D. Cloud Pub/Sub to Cloud SQL

Suggested Answer: *D*
Reference: https://cloud.google.com/pubsub/docs/ordering

Community vote distribution

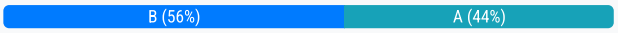

exampanic Highly Voted 6 years, 7 months ago

I believe the answer is B. "Pub/Sub doesn't provide guarantees about the order of message delivery. Strict message ordering can be achieved with buffering, often using Dataflow." https://cloud.google.com/solutions/data-lifecycle-cloud-platformthe correct is B https://cloud.google.com/pubsub/docs/stream-messages-dataflow"globally scaled frontend" means Dataflow is required. If the clients were in 1 region only pushing messages, then Pub/Sub alone would be the solution. Since there are global clients, this requires Dataflow.Now Pub/Sub guarantees message order. Until the exam does not change I would pick B.https://cloud.google.com/pubsub/docs/exactly-once-delivery referencePub/Sub supports exactly-once delivery, within a cloud region. The question states ¨global¨, so needs DataflowI believe that only the frontend is scaled globally. The backend API is the one that requires ordered delivery of the messages and guaranteed-once delivery of data. Currently, Pub/Sub supports ordered delivery within the same region (https://cloud.google.com/pubsub/docs/ordering#receiving_messages_in_order) and exactly-once delivery within the same region (https://cloud.google.com/pubsub/docs/exactly-once-delivery#exactly-once_delivery_guarantees). The right answer could be A, Pub/Sub alone.Answer is B. The question is talking about guaranteed-once FIFO delivery of data. Although Pub/sub provides data in order (FIFO) but it does 'at-least' once delivery of data. So, we need Dataflow for deduplication of data.I believe Pub/Sub now also supports exactly once delivery (in preview): https://cloud.google.com/pubsub/docs/exactly-once-delivery

upvoted 70 times

exampanic Highly Voted 6 years, 7 months ago

I believe the answer is B. "Pub/Sub doesn't provide guarantees about the order of message delivery. Strict message ordering can be achieved with buffering, often using Dataflow." https://cloud.google.com/solutions/data-lifecycle-cloud-platform

upvoted 70 times

zanfo 4 years, 5 months ago

the correct is B https://cloud.google.com/pubsub/docs/stream-messages-dataflow

upvoted 1 times

Clouddude123 1 year, 7 months ago

"globally scaled frontend" means Dataflow is required. If the clients were in 1 region only pushing messages, then Pub/Sub alone would be the solution. Since there are global clients, this requires Dataflow.

upvoted 2 times

TiagoM 5 years, 3 months ago

Now Pub/Sub guarantees message order. Until the exam does not change I would pick B.https://cloud.google.com/pubsub/docs/exactly-once-delivery referencePub/Sub supports exactly-once delivery, within a cloud region. The question states ¨global¨, so needs DataflowI believe that only the frontend is scaled globally. The backend API is the one that requires ordered delivery of the messages and guaranteed-once delivery of data. Currently, Pub/Sub supports ordered delivery within the same region (https://cloud.google.com/pubsub/docs/ordering#receiving_messages_in_order) and exactly-once delivery within the same region (https://cloud.google.com/pubsub/docs/exactly-once-delivery#exactly-once_delivery_guarantees). The right answer could be A, Pub/Sub alone.Answer is B. The question is talking about guaranteed-once FIFO delivery of data. Although Pub/sub provides data in order (FIFO) but it does 'at-least' once delivery of data. So, we need Dataflow for deduplication of data.I believe Pub/Sub now also supports exactly once delivery (in preview): https://cloud.google.com/pubsub/docs/exactly-once-delivery

upvoted 10 times

TiagoM 5 years, 3 months ago

Now Pub/Sub guarantees message order. Until the exam does not change I would pick B.

upvoted 10 times

melono 3 years, 10 months ago

https://cloud.google.com/pubsub/docs/exactly-once-delivery reference

upvoted 1 times

melono 3 years, 10 months ago

Pub/Sub supports exactly-once delivery, within a cloud region. The question states ¨global¨, so needs DataflowI believe that only the frontend is scaled globally. The backend API is the one that requires ordered delivery of the messages and guaranteed-once delivery of data. Currently, Pub/Sub supports ordered delivery within the same region (https://cloud.google.com/pubsub/docs/ordering#receiving_messages_in_order) and exactly-once delivery within the same region (https://cloud.google.com/pubsub/docs/exactly-once-delivery#exactly-once_delivery_guarantees). The right answer could be A, Pub/Sub alone.

upvoted 9 times

melono 3 years, 10 months ago

Pub/Sub supports exactly-once delivery, within a cloud region. The question states ¨global¨, so needs Dataflow

upvoted 9 times

CosminCiuc 3 years, 6 months ago

I believe that only the frontend is scaled globally. The backend API is the one that requires ordered delivery of the messages and guaranteed-once delivery of data. Currently, Pub/Sub supports ordered delivery within the same region (https://cloud.google.com/pubsub/docs/ordering#receiving_messages_in_order) and exactly-once delivery within the same region (https://cloud.google.com/pubsub/docs/exactly-once-delivery#exactly-once_delivery_guarantees). The right answer could be A, Pub/Sub alone.

upvoted 3 times

jask 4 years, 10 months ago

Answer is B. The question is talking about guaranteed-once FIFO delivery of data. Although Pub/sub provides data in order (FIFO) but it does 'at-least' once delivery of data. So, we need Dataflow for deduplication of data.I believe Pub/Sub now also supports exactly once delivery (in preview): https://cloud.google.com/pubsub/docs/exactly-once-delivery

upvoted 14 times

jask 4 years, 10 months ago

Answer is B. The question is talking about guaranteed-once FIFO delivery of data. Although Pub/sub provides data in order (FIFO) but it does 'at-least' once delivery of data. So, we need Dataflow for deduplication of data.

upvoted 14 times

emirhosseini 3 years, 10 months ago

I believe Pub/Sub now also supports exactly once delivery (in preview): https://cloud.google.com/pubsub/docs/exactly-once-delivery

upvoted 9 times

xhova 6 years, 3 months ago

B is the answer. CloudSQL is only for storage, to get the messages in order you need timestamp processed in dataflow to arrange them before putting it in any storage volume. The system described is not querying a db it is expecting a stream of messages only dataflow can correct the order. ACID has no value here because the db is not being queried. You'll not find any documentation on pub/sub order being corrected with a db. See notes below on pub/sub and dataflow using timestamps and windows to ensure order https://cloud.google.com/pubsub/docs/pubsub-dataflow

upvoted 28 times

chickenwingz Most Recent 1 year, 1 month ago

PubSub can only guarantee AT-LEAST-ONCE delivery, so dataflow is required afterwards to clean up any duplicates

upvoted 1 times

neokyle 1 year, 6 months ago

Answer is still B (Pub/Sub + Dataflow), because https://cloud.google.com/pubsub/docs/exactly-once-delivery#regional_considerations mentions "The exactly-once delivery guarantee only applies when subscribers connect to the service in the same region. If your subscriber application is spread across multiple regions, it can lead to duplicate message delivery, even when exactly-once delivery is enabled." (so the exactly-once delivery only works regionally, but the question says architect for global) The question's constraint of a globally scaled frontend, means you'll have clients in multiple regions.

upvoted 2 times

plumbig11 1 year, 7 months ago

Cloud Pub/Sub to Dataflow because the data is streaming.

upvoted 1 times

selected 1 year, 9 months ago

https://cloud.google.com/pubsub/docs/exactly-once-delivery#regional_considerations

upvoted 2 times

awsgcparch 2 years ago

Google Cloud Pub/Sub now supports message ordering, which ensures that messages with the same ordering key are delivered in the exact order they were published. This feature addresses the requirement for strict chronological order without the need for additional services. Key Features of Cloud Pub/Sub with Message Ordering: Message Ordering: By using ordering keys, Pub/Sub can guarantee that messages are delivered in the order they are published. Exactly-once Delivery: Pub/Sub supports at-least-once delivery and can be configured to handle duplicate messages. Scalability and Reliability: Pub/Sub is a fully managed service that scales automatically and ensures high availability.

upvoted 4 times

19040e5 2 years, 3 months ago

A. Cloud Pub/Sub alone Cloud Pub/Sub Ordered Delivery: Cloud Pub/Sub natively supports ordered delivery when using the same ordering key. This guarantees that messages with the same key are delivered to subscribers in the order they were published, preventing out-of-order events. Exactly-once Delivery: Cloud Pub/Sub also offers exactly-once delivery within a region, ensuring that each message is delivered to a subscriber only once.

upvoted 6 times

hitmax87 2 years, 3 months ago

Pub/Sub provide ordered delivery but doesn't ensure deduplication. Cloud SQL is regional resource and it cant perform custom logic. I go with B

upvoted 1 times

mesodan 2 years, 5 months ago

A is correct. While Cloud Dataflow can be used for data processing, it doesn't guarantee FIFO order on its own. Additionally, introducing another processing layer adds complexity and might not be necessary for this specific requirement.

upvoted 2 times

yas_cloud 2 years, 6 months ago

Dont see a need for SQL here as the answer suggests. Option B is more appropriate

upvoted 1 times

Devx198912233 2 years, 6 months ago

B would be correct as pubsub b service might redeliver messages. When you receive messages in order and the Pub/Sub service redelivers a message with an ordering key, Pub/Sub maintains order by also redelivering the subsequent messages with the same ordering key. The Pub/Sub service redelivers these messages in the order that it originally received them.

upvoted 2 times

theBestStudent 2 years, 10 months ago

Answer is B. Key words: globally scaled and ensure delivery only once: Pubsub + Dataflow. If it were only one region, it would be fine to say just pubsub, but it is globally scaled.

upvoted 1 times

Frusci 2 years, 11 months ago

B, you need dataflow to deduplicate. Pub/Sub does "at-least" once delivery.

upvoted 2 times

AL_everyday 3 years ago

It should be A

upvoted 1 times

jlambdan 3 years, 2 months ago

https://cloud.google.com/pubsub/docs/exactly-once-delivery

upvoted 2 times

Question #103 *Topic 1*

Your company is planning to perform a lift and shift migration of their Linux RHEL 6.5+ virtual machines. The virtual machines are running in an on-premises
VMware environment. You want to migrate them to Compute Engine following Google-recommended practices. What should you do?

A. 1. Define a migration plan based on the list of the applications and their dependencies. 2. Migrate all virtual machines into Compute Engine individually with Migrate for Compute Engine.

B. 1. Perform an assessment of virtual machines running in the current VMware environment. 2. Create images of all disks. Import disks on Compute Engine. 3. Create standard virtual machines where the boot disks are the ones you have imported.

C. 1. Perform an assessment of virtual machines running in the current VMware environment. 2. Define a migration plan, prepare a Migrate for Compute Engine migration RunBook, and execute the migration.

D. 1. Perform an assessment of virtual machines running in the current VMware environment. 2. Install a third-party agent on all selected virtual machines. 3. Migrate all virtual machines into Compute Engine.

Suggested Answer: *C*
The framework illustrated in the preceding diagram has four phases: ɹ€¢ Assess. In this phase, you assess your source environment, assess the workloads that you want to migrate to Google Cloud, and assess which VMs support each workload. ɹ€¢ Plan. In this phase, you create the basic infrastructure for Migrate for Compute Engine, such as provisioning the resource hierarchy and setting up network access. ɹ€¢ Deploy. In this phase, you migrate the VMs from the source environment to Compute Engine. ɹ€¢ Optimize. In this phase, you begin to take advantage of the cloud technologies and capabilities. Reference: https://cloud.google.com/architecture/migrating-vms-migrate-for-compute-engine-getting-started

Community vote distribution
C (93%)

kopper2019 Highly Voted 5 years, 1 month ago
Ans) C , Migrate for Compute Engine organizes groups of VMs into Waves. After understanding the dependencies of your applications, create runbooks that contain groups of VMs and begin your migration! https://cloud.google.com/migrate/compute-engine/docs/4.5/how-to/migrate-on-premises-to-gcp/overview
upvoted 38 times

technodev Highly Voted 4 years, 7 months ago
Selected Answer: C
I got this question in my exam.
upvoted 15 times

technodev 4 years, 7 months ago
I got this question in my exam.Did u passed...? If yes, then Congratulations and let me know the correct answer
upvoted 15 times

Sur_Nikki 3 years, 3 months ago
Did u passed...? If yes, then Congratulations and let me know the correct answer
upvoted 1 times

CtrlX Most Recent 11 months, 2 weeks ago
answer C. Was in my exam
upvoted 1 times

thatb0yjosh 1 year, 3 months ago
I chose C. The tool is also now called Migrate to Virtual Machines
upvoted 1 times

plumbig11 1 year, 7 months ago
1. Perform an assessment of virtual machines running in the current VMware environment. (You should as best pratice perform an assessment of virtual machines and define a migration plan) 2. Define a migration plan, prepare a Migrate for Compute Engine migration RunBook, and execute the migration. (compute Engine migration Runbook is the best pratice)
upvoted 1 times

3fd692e 1 year, 10 months ago

Assess, Plan, Migrate. Textbook perfect

upvoted 1 times

salim_ 3 years, 3 months ago

https://cloud.google.com/migrate/compute-engine/docs/4.11/how-to/migrate-on-premises-to-gcp/overview

upvoted 1 times

8d31d36 3 years, 6 months ago

The reason why Option B is preferable over Option C is that it involves creating images of all disks and importing them into Compute Engine, which can significantly reduce the amount of time required for the migration. Additionally, creating standard virtual machines from the imported disks is a straightforward process, and it ensures that the migrated virtual machines are identical to the on-premises virtual machines, which can simplify the migration process and minimize the risk of compatibility issues.

upvoted 2 times

examch 3 years, 7 months ago

C is the correct answer, Runbooks are created from the Migrate for Compute Engine Manager. The system queries VMware or AWS for VMs and generates a CSV for you to edit. By editing the CSV, you define: The VMs in a wave. The order in which those VMs are migrated. The type and disk space of VMs that are launched on Google Cloud. Other characteristics that are defined in the Runbook reference.
https://cloud.google.com/migrate/compute-engine/docs/4.8/how-to/organizing-migrations/creating-and-modifying-runbooks#generating_runbook_templates

upvoted 2 times

megumin 3 years, 9 months ago

C is ok

upvoted 1 times

AzureDP900 3 years, 10 months ago

C is most suitable for this use

upvoted 1 times

AzureDP900 4 years, 1 month ago

C is right, It defines all logical steps to migrate on-premise to google cloud.

upvoted 2 times

meokey 4 years, 4 months ago

Does Ans. C) still valid as of latest GCE 5.0? in the doc "Migrating VM groups" with version GCE 5.0, I do not see "runbook" anymore which is explained up to version GCE 4.8. https://cloud.google.com/migrate/compute-engine/docs/5.0/how-to/migrating-vm-groups

upvoted 2 times

gaojun 4 years, 4 months ago

Go for C

upvoted 1 times

haroldbenites 4 years, 8 months ago

Go for C.

upvoted 1 times

nikiwi 4 years, 8 months ago

why not A? seems pretty obvious if you look at the google doc: https://cloud.google.com/migrate/compute-engine/docs/5.0/concepts/lifecycleWhen you are doing cloud migrations, you do migrations in "waves" which are groupings of one or more applications/workloads. Moving machines individually would break things, such as dependencies. This is standard industry practice.

upvoted 1 times

nikiwi 4 years, 8 months ago

why not A? seems pretty obvious if you look at the google doc: https://cloud.google.com/migrate/compute-engine/docs/5.0/concepts/lifecycle

upvoted 1 times

atlasga 4 years, 7 months ago

When you are doing cloud migrations, you do migrations in "waves" which are groupings of one or more applications/workloads. Moving machines individually would break things, such as dependencies. This is standard industry practice.

upvoted 3 times

vincy2202 4 years, 8 months ago

C is the correct answer.

upvoted 1 times

Question #104 *Topic 1*

You need to deploy an application to Google Cloud. The application receives traffic via TCP and reads and writes data to the filesystem. The application does not support horizontal scaling. The application process requires full control over the data on the file system because concurrent access causes corruption. The business is willing to accept a downtime when an incident occurs, but the application must be available 24/7 to support their business operations. You need to design the architecture of this application on Google Cloud. What should you do?

A. Use a managed instance group with instances in multiple zones, use Cloud Filestore, and use an HTTP load balancer in front of the instances.

B. Use a managed instance group with instances in multiple zones, use Cloud Filestore, and use a network load balancer in front of the instances.

C. Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use an HTTP load balancer in front of the instances.

D. Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use a network load balancer in front of the instances.

Suggested Answer: *D*

Reference: https://cloud.google.com/compute/docs/instance-groups

Community vote distribution

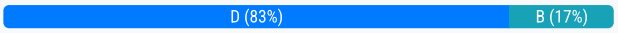

VishalB Highly Voted 4 years, 7 months ago

Correct Ans : D Since the Traffic is TCP, Ans A & C gets eliminated as HTTPS load balance is not supported. B - File storage system is Cloud Firestore which do not give full control, hence eliminated. D - Unmanaged instance group with network load balance with regional persistent disk for storage gives full control which is required for the migration.almost all good, except for File Storage, is not Cloud Firestore, it is a new service for sharing filesystems across VMs (like a NAS in a traditional infrastructure).what about the fact that is the unmanaged instance group is not regional , so you can't create it in more than 1 zone ?Can we group to running instances in different zones to an unmanaged instance group?

upvoted 64 times

VishalB Highly Voted 4 years, 7 months ago

Correct Ans : D Since the Traffic is TCP, Ans A & C gets eliminated as HTTPS load balance is not supported. B - File storage system is Cloud Firestore which do not give full control, hence eliminated. D - Unmanaged instance group with network load balance with regional persistent disk for storage gives full control which is required for the migration.

upvoted 64 times

poseidon24 4 years, 6 months ago

almost all good, except for File Storage, is not Cloud Firestore, it is a new service for sharing filesystems across VMs (like a NAS in a traditional infrastructure).

upvoted 10 times

kimharsh 3 years, 8 months ago

what about the fact that is the unmanaged instance group is not regional , so you can't create it in more than 1 zone ?Can we group to running instances in different zones to an unmanaged instance group?

upvoted 7 times

kimharsh 3 years, 8 months ago

what about the fact that is the unmanaged instance group is not regional , so you can't create it in more than 1 zone ?

upvoted 7 times

Jerryzzyy 2 years, 6 months ago

Can we group to running instances in different zones to an unmanaged instance group?

upvoted 1 times

kopper2019 4 years, 7 months ago

Ans) D , unmanaged instance group as application does not support horizontal scaling and network load balancer as no mention of http traffic .

upvoted 29 times

Pime13 6 months ago

Active-Standby (High Availability without Horizontal Scaling): Since the application does not support horizontal scaling and concurrent access causes corruption, you must ensure only one instance is active and writing to the disk at any given time. An unmanaged instance group is ideal for this manual "Active-Standby" configuration where you manage non-identical instances (one active, one standby) and control failover manually or via health checks.Regional Persistent Disk (Data Integrity & Zonal Resilience): A regional persistent disk (RPD) provides synchronous data replication across two zones in a single region. In this architecture, if the active instance in Zone A fails, the data is already available in Zone B. You can then "force-attach" the RPD to the standby instance in Zone B, ensuring the 24/7 availability requirement is met with minimal downtime. Network Load Balancer (TCP Traffic): The application receives TCP traffic. While HTTP load balancers are for Layer 7 traffic, a Network Load Balancer operates at Layer 4 and is designed to distribute TCP/UDP traffic. It can be configured with health checks to direct traffic to the currently active instance.Exclusive File Access: Managed Instance Groups (MIGs) are typically used for scaling identical clones. Even "Stateful MIGs" are generally intended for workloads like databases that handle their own replication. For a legacy app requiring strict "single-writer" access to prevent filesystem corruption, the manual control provided by an unmanaged group and RPD is the safest approach.

upvoted 1 times

Pime13 6 months ago

Selected Answer: D

Active-Standby (High Availability without Horizontal Scaling): Since the application does not support horizontal scaling and concurrent access causes corruption, you must ensure only one instance is active and writing to the disk at any given time. An unmanaged instance group is ideal for this manual "Active-Standby" configuration where you manage non-identical instances (one active, one standby) and control failover manually or via health checks.

upvoted 1 times

Pime13 6 months ago

Regional Persistent Disk (Data Integrity & Zonal Resilience): A regional persistent disk (RPD) provides synchronous data replication across two zones in a single region. In this architecture, if the active instance in Zone A fails, the data is already available in Zone B. You can then "force-attach" the RPD to the standby instance in Zone B, ensuring the 24/7 availability requirement is met with minimal downtime. Network Load Balancer (TCP Traffic): The application receives TCP traffic. While HTTP load balancers are for Layer 7 traffic, a Network Load Balancer operates at Layer 4 and is designed to distribute TCP/UDP traffic. It can be configured with health checks to direct traffic to the currently active instance.Exclusive File Access: Managed Instance Groups (MIGs) are typically used for scaling identical clones. Even "Stateful MIGs" are generally intended for workloads like databases that handle their own replication. For a legacy app requiring strict "single-writer" access to prevent filesystem corruption, the manual control provided by an unmanaged group and RPD is the safest approach.

upvoted 1 times

Pime13 6 months ago

Regional Persistent Disk (Data Integrity & Zonal Resilience): A regional persistent disk (RPD) provides synchronous data replication across two zones in a single region. In this architecture, if the active instance in Zone A fails, the data is already available in Zone B. You can then "force-attach" the RPD to the standby instance in Zone B, ensuring the 24/7 availability requirement is met with minimal downtime.

Network Load Balancer (TCP Traffic): The application receives TCP traffic. While HTTP load balancers are for Layer 7 traffic, a Network Load Balancer operates at Layer 4 and is designed to distribute TCP/UDP traffic. It can be configured with health checks to direct traffic to the currently active instance.

upvoted 1 times

Pime13 Most Recent 6 months ago

Exclusive File Access: Managed Instance Groups (MIGs) are typically used for scaling identical clones. Even "Stateful MIGs" are generally intended for workloads like databases that handle their own replication. For a legacy app requiring strict "single-writer" access to prevent filesystem corruption, the manual control provided by an unmanaged group and RPD is the safest approach.

upvoted 1 times

Polosaty 1 year, 11 months ago

In Unmanaged Instance Group instances cannot be in different zones. I think that correct is D but maybe a mistake in the question.

upvoted 1 times

OrangeTiger 2 years ago

Why not a B? Because the application doesn't support for horizonal scale. I chose D.

upvoted 1 times

duzapo 2 years, 5 months ago

D is corrrect TCP load balancer plus UNMANAGED

upvoted 2 times

TheCloudGuruu 2 years, 9 months ago

must be unmanaged

upvoted 1 times

JC0926 2 years, 11 months ago

Since the application does not support horizontal scaling, a managed instance group is not required. Instead, an unmanaged instance group can be used to ensure that the application runs on multiple instances in different zones for high availability. The network load balancer is designed to handle TCP and UDP traffic The HTTP(S) load balancer is designed specifically for HTTP and HTTPS traffic.Thanks for the apt explanation

upvoted 18 times

JC0926 2 years, 11 months ago

Selected Answer: D

Since the application does not support horizontal scaling, a managed instance group is not required. Instead, an unmanaged instance group can be used to ensure that the application runs on multiple instances in different zones for high availability. The network load balancer is designed to handle TCP and UDP traffic The HTTP(S) load balancer is designed specifically for HTTP and HTTPS traffic.

upvoted 18 times

Sur_Nikki 2 years, 9 months ago

Thanks for the apt explanation

upvoted 2 times

omermahgoub 3 years, 1 month ago

An unmanaged instance group allows you to create and manage a group of Compute Engine instances manually, rather than using an autoscaling solution like a managed instance group. This is appropriate for an application that does not support horizontal scaling, as you can manually create and manage the number of instances needed to meet the traffic demands.

To ensure high availability and minimize downtime, you should deploy the instances in different zones and use a regional persistent disk to store the application's data. This will ensure that the application is still available even if one of the instances or a zone experiences an outage.

A network load balancer should be used in front of the instances to distribute traffic to the instances. A network load balancer is a highly available and scalable load balancing solution that operates at the network layer and can handle high volumes of traffic. It can also balance traffic across multiple zones to ensure that the application is always available to users.

upvoted 22 times

omermahgoub 3 years, 1 month ago

An unmanaged instance group allows you to create and manage a group of Compute Engine instances manually, rather than using an autoscaling solution like a managed instance group. This is appropriate for an application that does not support horizontal scaling, as you can manually create and manage the number of instances needed to meet the traffic demands. To ensure high availability and minimize downtime, you should deploy the instances in different zones and use a regional persistent disk to store the application's data. This will ensure that the application is still available even if one of the instances or a zone experiences an outage. A network load balancer should be used in front of the instances to distribute traffic to the instances. A network load balancer is a highly available and scalable load balancing solution that operates at the network layer and can handle high volumes of traffic. It can also balance traffic across multiple zones to ensure that the application is always available to users.Therefore, the correct answer is option D: Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use a network load balancer in front of the instances.

upvoted 22 times

omermahgoub 3 years, 1 month ago

Therefore, the correct answer is option D: Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use a network load balancer in front of the instances.

upvoted 10 times

oms_muc 3 years, 1 month ago

Regional Persistent Disk, as App requires full control of filesystem data without concurrent access (block storage vs. file storage (NAS). https://cloud.google.com/compute/docs/instance-groups https://cloud.google.com/load-balancing/docs/choosing-load-balancer

upvoted 1 times

surajkrishnamurthy 3 years, 2 months ago

Option A & B eliminated because we cannot use managed instance group since the app does not support Horizontal scaling Option C > HTTP load balancer is Layer 7 & application is receiving traffic via TCP Option D > is best answer because we are using Network load balancer Layer 4 which meets the condition "application receives traffic via TCP"

upvoted 5 times

fiercedog 3 years, 2 months ago

Checking the comparative analysis of storage options, we can see that Filestore is not suitable for the workload, hence A and B are out. C is out because it restricts to HTTP traffic. Answer is Dhttps://cloud.google.com/architecture/storage-advisor#comparative_analysis

upvoted 2 times

fiercedog 3 years, 2 months ago

Selected Answer: D

Checking the comparative analysis of storage options, we can see that Filestore is not suitable for the workload, hence A and B are out. C is out because it restricts to HTTP traffic. Answer is D

upvoted 2 times

fiercedog 3 years, 2 months ago

https://cloud.google.com/architecture/storage-advisor#comparative_analysis

upvoted 1 times

megumin 3 years, 3 months ago

D is ok

upvoted 1 times

SerGCP 3 years, 3 months ago

D is not possibile you cannot create a regional unmanaged instance group. https://cloud.google.com/compute/docs/instance-groups/creating-groups-of-unmanaged-instancesDowntime is acceptable, Disk is regional, in case of issue the unmanaged instance group can be moved to other zone, disk has data. A&B are not at all an option

upvoted 3 times

SerGCP 3 years, 3 months ago

Selected Answer: B

D is not possibile you cannot create a regional unmanaged instance group. https://cloud.google.com/compute/docs/instance-groups/creating-groups-of-unmanaged-instances

upvoted 3 times

KyubiBlaze 3 years, 2 months ago

Downtime is acceptable, Disk is regional, in case of issue the unmanaged instance group can be moved to other zone, disk has data. A&B are not at all an option

upvoted 3 times

minmin2020 3 years, 4 months ago

D. Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use a network load balancer in front of the instances. D is the only option as the application does not support horizontal scaling (no MIG), it needs full control (no filestore) and has TCP traffic (no HTTP LB).

upvoted 1 times

Santanu_01 3 years, 4 months ago

D is more appropriate solution ---In the ques. it is mentioned done not support horizontal scaling --> hence Unmanaged Instance , and traffic is TCP --- hence N/W Load balancer

upvoted 1 times

Rajeev26 3 years, 4 months ago

does not support horizontal scaling so MIG not needed. TCP traffic so NW LB is ok

upvoted 1 times

Question #105 *Topic 1*

Your company has an application running on multiple Compute Engine instances. You need to ensure that the application can communicate with an on-premises service that requires high throughput via internal IPs, while minimizing latency. What should you do?

A. Use OpenVPN to configure a VPN tunnel between the on-premises environment and Google Cloud.

B. Configure a direct peering connection between the on-premises environment and Google Cloud.

C. Use Cloud VPN to configure a VPN tunnel between the on-premises environment and Google Cloud.

D. Configure a Cloud Dedicated Interconnect connection between the on-premises environment and Google Cloud.

Suggested Answer: *C*

Reference - https://cloud.google.com/architecture/setting-up-private-access-to-cloud-apis-through-vpn-tunnels

Community vote distribution

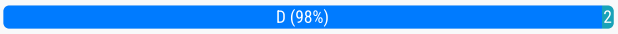

kopper2019 Highly Voted 5 years, 1 month ago

Ans) D , Reason : high throughput via internal IPsThis is tricky questions , it can be achieved by C and D ... Multiple Computes and Costs .. they are trying to test knowledge on VPN and Tunnels ....

upvoted 67 times

kopper2019 Highly Voted 5 years, 1 month ago

Ans) D , Reason : high throughput via internal IPs

upvoted 67 times

ShadowLord 4 years ago

This is tricky questions , it can be achieved by C and D ... Multiple Computes and Costs .. they are trying to test knowledge on VPN and Tunnels ....

upvoted 2 times

XDevX 5 years, 1 month ago

IMHO the correct answer is D. Reason: "requires high throughput via internal IPs, while minimizing latency" - both are aspects you cannot guarantee with using VPN traversing the internet.

upvoted 23 times

the1dv 2 years, 7 months ago

Should be D as high throughput

upvoted 1 times

Gilbaliano 2 years, 8 months ago

Justo be DShould be D

upvoted 1 times

Gilbaliano 2 years, 8 months ago

Selected Answer: D

Justo be D

upvoted 1 times

Gilbaliano 2 years, 8 months ago

Should be D

upvoted 1 times

kamradamir 2 years, 8 months ago

It should be D, since need to communicate via Internal IP

upvoted 1 times

thewalker 2 years, 9 months ago

Selected Answer: C

Communication through internal IPs - VPN. So, C

upvoted 1 times

thewalker 2 years, 9 months ago

Communication through internal IPs - VPN. So, CNO, VPN can't meet the low latency and high throughput requirements. dedicated interconnect is a connection between you and GCP VPC. So, you can use any IP address range (private RFC 1918 or public IP ranges)yes, you are very smart but

question says high throughput and minimizing latency > Interconnect

upvoted 1 times

halifax Most Recent 1 year, 5 months ago

NO, VPN can't meet the low latency and high throughput requirements. dedicated interconnect is a connection between you and GCP VPC. So, you can use any IP address range (private RFC 1918 or public IP ranges)

upvoted 1 times

ccpmad 2 years, 2 months ago

yes, you are very smart but question says high throughput and minimizing latency > Interconnect

upvoted 2 times

duzapo 2 years, 11 months ago

D cause you need high throughput

upvoted 2 times

marcjimz 3 years ago

D - high throughput

upvoted 1 times

wooloo 3 years ago

How is C marked a correct one?

upvoted 2 times

Danomine416 3 years, 3 months ago

Initially thought 'D' but the question says 'high throughput via internal IPs' so go with VPN answer 'C'

upvoted 2 times

grejao 3 years, 4 months ago

It can´t be VPN, Only interconnect can minimizing latency. D is the right answer.

upvoted 1 times

zerg0 3 years, 6 months ago

Internal IP + high throughput

upvoted 5 times

WFCheong 3 years, 7 months ago

Ans) D , Reason : high throughput via internal IPs

upvoted 3 times

omermahgoub 3 years, 8 months ago

D. Configure a Cloud Dedicated Interconnect connection between the on-premises environment and Google Cloud.

A Cloud Dedicated Interconnect is a high-bandwidth, low-latency network connection that allows you to connect your on-premises environment to Google Cloud Platform (GCP) using a dedicated network connection. It provides a direct physical connection between your on-premises network and GCP, which can help to reduce latency and increase the throughput of your application.

upvoted 4 times

omermahgoub 3 years, 8 months ago

D. Configure a Cloud Dedicated Interconnect connection between the on-premises environment and Google Cloud. A Cloud Dedicated Interconnect is a high-bandwidth, low-latency network connection that allows you to connect your on-premises environment to Google Cloud Platform (GCP) using a dedicated network connection. It provides a direct physical connection between your on-premises network and GCP, which can help to reduce latency and increase the throughput of your application.Option A, using OpenVPN to configure a VPN tunnel between the on-premises environment and Google Cloud, is not a recommended approach. OpenVPN is a free and open-source software application that implements virtual private network (VPN) techniques to create secure point-to-point connections. While OpenVPN can be used to establish a VPN tunnel between an on-premises environment and GCP, it may not provide the level of performance and reliability required for high-throughput applications. Option B, configuring a direct peering connection between the on-premises environment and Google Cloud, is not a recommended approach. Direct Peering is a high-bandwidth, low-latency network connection that allows you to connect your on-premises network directly to Google's network. While Direct Peering can be used to connect your on-premises environment to GCP, it is typically used for high-bandwidth workloads such as video streaming and may not be suitable for applications that require low latency. Option C, using Cloud VPN to configure a VPN tunnel between the on-premisesCloud VPN, would also involve routing traffic over the public internet and would not provide the low latency and high throughput that you need.Yeah, direct peering is high bandwidth and low latency, but may not be suitable for applications that require low latency. Good job, ChatGPT...

upvoted 4 times

omermahgoub 3 years, 8 months ago

Option A, using OpenVPN to configure a VPN tunnel between the on-premises environment and Google Cloud, is not a recommended approach. OpenVPN is a free and open-source software application that implements virtual private network (VPN) techniques to create secure point-to-point

connections. While OpenVPN can be used to establish a VPN tunnel between an on-premises environment and GCP, it may not provide the level of performance and reliability required for high-throughput applications. Option B, configuring a direct peering connection between the on-premises environment and Google Cloud, is not a recommended approach. Direct Peering is a high-bandwidth, low-latency network connection that allows you to connect your on-premises network directly to Google's network. While Direct Peering can be used to connect your on-premises environment to GCP, it is typically used for high-bandwidth workloads such as video streaming and may not be suitable for applications that require low latency. Option C, using Cloud VPN to configure a VPN tunnel between the on-premisesCloud VPN, would also involve routing traffic over the public internet and would not provide the low latency and high throughput that you need.Yeah, direct peering is high bandwidth and low latency, but may not be suitable for applications that require low latency. Good job, ChatGPT...

upvoted 2 times

omermahgoub 3 years, 8 months ago

Option A, using OpenVPN to configure a VPN tunnel between the on-premises environment and Google Cloud, is not a recommended approach. OpenVPN is a free and open-source software application that implements virtual private network (VPN) techniques to create secure point-to-point connections. While OpenVPN can be used to establish a VPN tunnel between an on-premises environment and GCP, it may not provide the level of performance and reliability required for high-throughput applications.

Option B, configuring a direct peering connection between the on-premises environment and Google Cloud, is not a recommended approach. Direct Peering is a high-bandwidth, low-latency network connection that allows you to connect your on-premises network directly to Google's network. While Direct Peering can be used to connect your on-premises environment to GCP, it is typically used for high-bandwidth workloads such as video streaming and may not be suitable for applications that require low latency.

Option C, using Cloud VPN to configure a VPN tunnel between the on-premises

upvoted 2 times

omermahgoub 3 years, 8 months ago

Cloud VPN, would also involve routing traffic over the public internet and would not provide the low latency and high throughput that you need.

upvoted 1 times

medi01 3 years, 3 months ago

Yeah, direct peering is high bandwidth and low latency, but may not be suitable for applications that require low latency. Good job, ChatGPT...

upvoted 2 times

SureshbabuK 3 years, 8 months ago

Highthoughput connection VPN - No Interconnect - Y via Internal IP addresses. VPN - yes Interconnect - Y Interconnect yes for both so D

upvoted 3 times

megumin 3 years, 9 months ago

D is ok

upvoted 1 times

przema86 3 years, 9 months ago

high throughput => D

upvoted 1 times

Question #106 *Topic 1*

You are managing an application deployed on Cloud Run for Anthos, and you need to define a strategy for deploying new versions of the application. You want to evaluate the new code with a subset of production traffic to decide whether to proceed with the rollout. What should you do?

A. Deploy a new revision to Cloud Run with the new version. Configure traffic percentage between revisions.

B. Deploy a new service to Cloud Run with the new version. Add a Cloud Load Balancing instance in front of both services.

C. In the Google Cloud Console page for Cloud Run, set up continuous deployment using Cloud Build for the development branch. As part of the Cloud Build trigger, configure the substitution variable TRAFFIC_PERCENTAGE with the percentage of traffic you want directed to a new version.

D. In the Google Cloud Console, configure Traffic Director with a new Service that points to the new version of the application on Cloud Run. Configure Traffic Director to send a small percentage of traffic to the new version of the application.

Suggested Answer: *C*
Official Answer: C

Community vote distribution
A (74%) C (26%)

VishalB Highly Voted 4 years, 7 months ago
 Correct Answer: A o Each deployment to a service creates a revision. A revision consists of a specific container image, along with environment settings such as environment variables, memory limits, or concurrency value. o Once the new revision is deployed to a Service you can manage the traffic using MANAGE TRAFFIC option inside the revision tab https://cloud.google.com/run/docs/resource-model
upvoted 55 times

omermahgoub Highly Voted 3 years, 1 month ago
The correct answer is A. Deploy a new revision to Cloud Run with the new version. Configure traffic percentage between revisions.

Cloud Run for Anthos allows you to deploy new revisions of your application with a specific percentage of traffic, which allows you to perform a gradual rollout of the new version. To do this, you can follow these steps:

Deploy a new revision of your application to Cloud Run with the new version.

In the Cloud Run for Anthos console, navigate to the service that you want to roll out the new version for.

In the "Revisions" tab, you should see the new revision listed alongside the current revision.

Use the traffic percentage slider to specify the percentage of traffic that you want to send to the new revision. You can set the percentage to a small value initially, such as 5%, and gradually increase it over time as you evaluate the new version.

Once you have set the traffic percentage, Cloud Run for Anthos will start directing a portion of the traffic to the new revision, allowing you to evaluate the new version with a subset of production traffic.
upvoted 20 times

omermahgoub 3 years, 1 month ago
The correct answer is A. Deploy a new revision to Cloud Run with the new version. Configure traffic percentage between revisions. Cloud Run for Anthos allows you to deploy new revisions of your application with a specific percentage of traffic, which allows you to perform a gradual rollout of the new version. To do this, you can follow these steps: Deploy a new revision of your application to Cloud Run with the new version. In the Cloud Run for Anthos console, navigate to the service that you want to roll out the new version for. In the "Revisions" tab, you should see the new revision listed alongside the current revision. Use the traffic percentage slider to specify the percentage of traffic that you want to send to the new revision. You can set the percentage to a small value initially, such as 5%, and gradually increase it over time as you evaluate the new version. Once you have set the traffic percentage, Cloud Run for Anthos will start directing a portion of the traffic to the new revision, allowing you to evaluate the new version with a subset of production traffic.Extremely convinced by your explanations..Have u given this exam?Those Answers are surely generated by ChatgGPTYes, chatgpt did the exam, LOLOption B, deploying a new service and adding a Cloud Load Balancer instance in front of both services, is not recommended

because it would require you to create and manage a separate service for the new version, which would be more complex and less efficient than deploying a new revision. Option C, using continuous deployment with Cloud Build and substitution variables, is not relevant to this scenario because it involves deploying new versions automatically based on changes to a development branch, rather than manually deploying new revisions with a specific percentage of traffic. Option D, using Traffic Director, is also not relevant because Traffic Director is used for managing traffic between different services or clusters, rather than between revisions of

upvoted 20 times

Sur_Nikki 2 years, 9 months ago

Extremely convinced by your explanations..Have u given this exam?

upvoted 1 times

Sur_Nikki 2 years, 9 months ago

Extremely convinced by your explanations..Have u given this exam?Those Answers are surely generated by ChatgGPTYes, chatgpt did the exam, LOL

upvoted 1 times

baertierchen 2 years, 8 months ago

Those Answers are surely generated by ChatgGPT

upvoted 1 times

ccpmad 1 year, 8 months ago

Yes, chatgpt did the exam, LOL

upvoted 2 times

omermahgoub 3 years, 1 month ago

Option B, deploying a new service and adding a Cloud Load Balancer instance in front of both services, is not recommended because it would require you to create and manage a separate service for the new version, which would be more complex and less efficient than deploying a new revision. Option C, using continuous deployment with Cloud Build and substitution variables, is not relevant to this scenario because it involves deploying new versions automatically based on changes to a development branch, rather than manually deploying new revisions with a specific percentage of traffic. Option D, using Traffic Director, is also not relevant because Traffic Director is used for managing traffic between different services or clusters, rather than between revisions of the same service.

upvoted 12 times

CtrlX Most Recent 11 months, 2 weeks ago

answer A. Was in my exam

upvoted 1 times

plumbig11 1 year, 7 months ago

Deploy a new revision to Cloud Run with the new version. Configure traffic percentage between revisions.

upvoted 2 times

02fc23a 2 years, 2 months ago

A: https://cloud.google.com/run/docs/rollouts-rollbacks-traffic-migration#gradual

upvoted 1 times

odacir 2 years, 3 months ago

La Opción A - El uso de dashboards predefinidos proporciona una visión inmediata y eficiente del estado del sistema, y la capacidad de agregar métricas personalizadas y crear políticas de alertas permite una respuesta rápida y efectiva a los incidentes. Generar un Dashboard por incidente introduce complejidad innecesaria.

upvoted 1 times

odacir 2 years, 3 months ago

https://cloud.google.com/anthos/run/docs/rollouts-rollbacks-traffic-migration

upvoted 1 times

JPA210 2 years, 4 months ago

You all forget that in the exam they expect you to choose the most complete answer, and C is the most efficient and complete one, where you show knowledge of Cloud Build also, and in the real life that is how you should implement this solution.

upvoted 1 times

JPA210 2 years, 4 months ago

You all forget that in the exam they expect you to choose the most complete answer, and C is the most efficient and complete one, where you show knowledge of Cloud Build also, and in the real life that is how you should implement this solution.C is saying "from development branch"!

upvoted 1 times

SSS987 2 years, 1 month ago

C is saying "from development branch"!

upvoted 2 times

HRS1954 2 years, 5 months ago

A is correct - https://cloud.google.com/anthos/run/docs/rollouts-rollbacks-traffic-migration#gradual

upvoted 2 times

BigfootPanda 2 years, 7 months ago

C is ok as we want to DEFINE A STRATEGY

upvoted 3 times

TheCloudGuruu 2 years, 9 months ago

Answer is A

upvoted 1 times

zerg0 3 years ago

https://cloud.google.com/anthos/run/docs/rollouts-rollbacks-traffic-migration

upvoted 4 times

KM0107 3 years, 2 months ago

Selected A

upvoted 1 times

megumin 3 years, 3 months ago

A is ok

upvoted 1 times

sfsdeniso 3 years, 3 months ago

A is correct currently there is possibility to use tags to test in production without receiving real traffic:
https://cloud.google.com/anthos/run/docs/rollouts-rollbacks-traffic-migration#tags

upvoted 2 times

sfsdeniso 3 years, 4 months ago

whats wrong with C is configuring deployment from 'development branch' this is supper ugly

upvoted 2 times

zr79 3 years, 4 months ago

you can do a lab on this >>>Deploy Your Website on Cloud Run<<< Manage traffic is there

upvoted 1 times

Question #107 *Topic 1*

You are monitoring Google Kubernetes Engine (GKE) clusters in a Cloud Monitoring workspace. As a Site Reliability Engineer (SRE), you need to triage incidents quickly. What should you do?

A. Navigate the predefined dashboards in the Cloud Monitoring workspace, and then add metrics and create alert policies.

B. Navigate the predefined dashboards in the Cloud Monitoring workspace, create custom metrics, and install alerting software on a Compute Engine instance.

C. Write a shell script that gathers metrics from GKE nodes, publish these metrics to a Pub/Sub topic, export the data to BigQuery, and make a Data Studio dashboard.

D. Create a custom dashboard in the Cloud Monitoring workspace for each incident, and then add metrics and create alert policies.

Suggested Answer: *D*

Reference: https://cloud.google.com/monitoring/charts/dashboards

Community vote distribution

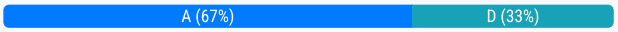

kopper2019 Highly Voted 4 years, 7 months ago

Ans) A .

upvoted 57 times

DiegoMDZ Highly Voted 4 years, 7 months ago

It's A for me... Create a dashboard for each incident?? I think D isn't a good choice...

upvoted 31 times

DiegoMDZ 4 years, 7 months ago

It's A for me... Create a dashboard for each incident?? I think D isn't a good choice...Yeah, creating a new dashboard for each incident doesn't seem like the quickest option.

upvoted 31 times

bandegg 2 years, 1 month ago

Yeah, creating a new dashboard for each incident doesn't seem like the quickest option.

upvoted 3 times

Pime13 Most Recent 6 months ago

Speed and Efficiency: Predefined GKE dashboards are "Kubernetes-native" and provide an immediate, comprehensive view of cluster health (CPU, memory, storage) and resource associations (nodes, pods, containers) without manual setup. Integrated Troubleshooting: These dashboards allow you to filter data by specific resources and transition contextually between logs, metrics, and incident details to find root causes in minutes. Proactive Monitoring: Adding relevant metrics and alert policies to these existing views follows SRE best practices by ensuring urgent, actionable conditions are detected and reported through standard notification channels.

upvoted 1 times

Mikeliz 1 year, 5 months ago

A is a better option, you don't need to create a dashboard for each incident

upvoted 3 times

plumbig11 1 year, 7 months ago

In this case predefined with metrics is enought

upvoted 3 times

6a8c7ad 1 year, 8 months ago

quickly would mean custom, and you can call it what you want. D

upvoted 1 times

hitmax87 1 year, 9 months ago

Selected Answer: D

If you need extended functionality you should create your dashboards, metrics and alerts. Dont mess existing ones.

upvoted 1 times

hitmax87 1 year, 9 months ago

If you need extended functionality you should create your dashboards, metrics and alerts. Dont mess existing ones.You would never create a new dashboard for each incident, which is what D is proposing.

upvoted 1 times

ryaryarya 1 year, 7 months ago

You would never create a new dashboard for each incident, which is what D is proposing.

upvoted 1 times

coolie1234 1 year, 9 months ago

Key is custom dashboard

upvoted 1 times

Gino17m 1 year, 10 months ago

A is correct

upvoted 3 times

mesodan 1 year, 11 months ago

A is correct. Option D is highly inefficient and time-consuming. Creating individual dashboards for every incident is impractical and slows down the triage process.

upvoted 5 times

OrangeTiger 2 years ago

I will go with A. uhhm, opinions are divided. In such cases, Q is often not good.

upvoted 1 times

hzaoui 2 years ago

Explanation: Cloud Monitoring provides predefined dashboards for monitoring GKE clusters, which facilitate an immediate and comprehensive view of cluster performance and health. As an SRE, utilizing these dashboards helps triage incidents quickly. You can also add additional metrics that are pertinent to the incident and create alert policies that will notify you when specific conditions indicative of an incident are met. This strategy allows for the proactive monitoring of incidents and rapid response when necessary.

upvoted 7 times

SSS987 2 years, 1 month ago

Ans: D. Although creating dashboard per incident sounds confusing and inefficient, it is still better than the impossible option A as we can't edit or add metrics to a predefined dashboard. Inefficient option vs Impossible Option - Inefficient one is ok!

upvoted 3 times

SSS987 2 years, 1 month ago

Ans: D. Although creating dashboard per incident sounds confusing and inefficient, it is still better than the impossible option A as we can't edit or add metrics to a predefined dashboard. Inefficient option vs Impossible Option - Inefficient one is ok!Optipn A is possible. You can't add widgets to predefined dashboard but you can create alert policies based on metrics. Although the structure of the question is actually misleading

upvoted 3 times

Gino17m 1 year, 10 months ago

Optipn A is possible. You can't add widgets to predefined dashboard but you can create alert policies based on metrics. Although the structure of the question is actually misleading

upvoted 2 times

ade7cae 2 years, 2 months ago

A is correct. For D, creating dashboards for each incident isn't practical

upvoted 3 times

spuyol 2 years, 2 months ago

Currently you can not modify pre-defined dashboards. Has no sense to create a dashboard for each incident. Conclusion: no answer is good enough

upvoted 1 times

brentc 2 years, 3 months ago

We do have dashboards for each incident

upvoted 2 times

TopTalk 2 years, 4 months ago

"You can't delete or modify the automatically-created dashboards; however, when support for copying the dashboard exists, you can modify the copy. In general, you can also copy charts on a predefined dashboard to a dashboard that you create. Dashboards that you create are custom dashboards. Custom dashboards let you display information that is of interest to you, organized in a way that's useful to you. "
https://cloud.google.com/monitoring/charts/predefined-dashboardsCreate a dashboard for each inc ????? serious

upvoted 3 times

TopTalk 2 years, 4 months ago

Selected Answer: D

"You can't delete or modify the automatically-created dashboards; however, when support for copying the dashboard exists, you can modify the copy. In general, you can also copy charts on a predefined dashboard to a dashboard that you create. Dashboards that you create are custom dashboards. Custom dashboards let you display information that is of interest to you, organized in a way that's useful to you. "
https://cloud.google.com/monitoring/charts/predefined-dashboards

upvoted 3 times

ArtistS 2 years, 3 months ago

Create a dashboard for each inc ????? serious

upvoted 1 times

Question #108 *Topic 1*

You are implementing a single Cloud SQL MySQL second-generation database that contains business-critical transaction data. You want to ensure that the minimum amount of data is lost in case of catastrophic failure. Which two features should you implement? (Choose two.)

A. Sharding

B. Read replicas

C. Binary logging

D. Automated backups

E. Semisynchronous replication

Suggested Answer: *C, D*

Backups help you restore lost data to your Cloud SQL instance. Additionally, if an instance is having a problem, you can restore it to a previous state by using the backup to overwrite it. Enable automated backups for any instance that contains necessary data. Backups protect your data from loss or damage. Enabling automated backups, along with binary logging, is also required for some operations, such as clone and replica creation. Reference: https://cloud.google.com/sql/docs/mysql/backup-recovery/backups

Community vote distribution



kopper2019 Highly Voted 4 years, 7 months ago

Ans) C and D Cloud SQL. If you use Cloud SQL, the fully managed Google Cloud MySQL database, you should enable automated backups and binary logging for your Cloud SQL instances. This allows you to perform a point-in-time recovery, which restores your database from a backup and recovers it to a fresh Cloud SQL instanceAnd: a read-replica won't help against "catastrophic failures" like accidental deletionscatastrophic failure means disaster like a zonal datacenter level failure or regional failure

upvoted 40 times

kopper2019 Highly Voted 4 years, 7 months ago

Ans) C and D Cloud SQL. If you use Cloud SQL, the fully managed Google Cloud MySQL database, you should enable automated backups and binary logging for your Cloud SQL instances. This allows you to perform a point-in-time recovery, which restores your database from a backup and recovers it to a fresh Cloud SQL instance

upvoted 40 times

HenkH 3 years, 3 months ago

And: a read-replica won't help against "catastrophic failures" like accidental deletionscatastrophic failure means disaster like a zonal datacenter level failure or regional failure

upvoted 5 times

HenkH 3 years, 3 months ago

And: a read-replica won't help against "catastrophic failures" like accidental deletions

upvoted 5 times

RVivek 3 years ago

catastrophic failure means disaster like a zonal datacenter level failure or regional failure

upvoted 4 times

victory108 4 years, 7 months ago

C. Binary logging D. Automated backups

upvoted 11 times

Pime13 Most Recent 6 months ago

Selected Answer: CD

Binary Logging (C): This is the most critical feature for achieving a low Recovery Point Objective (RPO). Binary logs record all data changes (insertions, updates, deletions) made to the database. In the event of a failure, these logs allow you to perform point-in-time recovery, restoring your data to the exact state it was in just moments before the catastrophe occurred. Automated Backups (D): This feature ensures that a complete, consistent copy of your database is created every day without manual intervention. These backups serve as the foundational "base" for any restoration process. When combined with binary logging, you can restore a daily backup and then "play back" the binary logs to recover all subsequent transactions.

upvoted 1 times

Pime13 6 months ago

Binary Logging (C): This is the most critical feature for achieving a low Recovery Point Objective (RPO). Binary logs record all data changes (insertions, updates, deletions) made to the database. In the event of a failure, these logs allow you to perform point-in-time recovery, restoring your data to the exact state it was in just moments before the catastrophe occurred. Automated Backups (D): This feature ensures that a complete, consistent copy of your database is created every day without manual intervention. These backups serve as the foundational "base" for any restoration process. When combined with binary logging, you can restore a daily backup and then "play back" the binary logs to recover all subsequent transactions.Option A (Sharding): This is a horizontal scaling technique used to distribute data across multiple instances to improve performance. It is not a data protection or disaster recovery mechanism. Option B (Read Replicas): Replicas are primarily used to scale read-heavy workloads. While they can sometimes be promoted to a primary instance, they are not a substitute for backups and point-in-time recovery, as accidental data deletions on the primary are instantly replicated to the replicas. Option E (Semisynchronous Replication): While this improves data consistency between a primary and a replica, it is a high-availability feature rather than a recovery feature. It does not protect against data corruption or the need to restore from an earlier point in time.

upvoted 1 times

Pime13 6 months ago

Option A (Sharding): This is a horizontal scaling technique used to distribute data across multiple instances to improve performance. It is not a data protection or disaster recovery mechanism. Option B (Read Replicas): Replicas are primarily used to scale read-heavy workloads. While they can sometimes be promoted to a primary instance, they are not a substitute for backups and point-in-time recovery, as accidental data deletions on the primary are instantly replicated to the replicas. Option E (Semisynchronous Replication): While this improves data consistency between a primary and a replica, it is a high-availability feature rather than a recovery feature. It does not protect against data corruption or the need to restore from an earlier point in time.

upvoted 1 times

plumbig11 1 year, 7 months ago

Binary logging and automated backups

upvoted 1 times

odacir 2 years, 3 months ago

Prerequisites for creating a read replica Before you can create a read replica of a primary Cloud SQL instance, the instance must meet the following requirements: Automated backups must be enabled. Binary logging must be enabled which requires point-in-time recovery to be enabled. Learn more about the impact of these logs. At least one backup must have been created after binary logging was enabled.
https://cloud.google.com/sql/docs/mysql/replication#requirements

upvoted 6 times

someone2011 2 years, 5 months ago

CD Before being able to create a read replica, you have to make sure "binary logging and automated backup" are enabled. So picking only D or C without the other one makes no sense. https://cloud.google.com/sql/docs/mysql/replication/create-replica

upvoted 1 times

HRS1954 2 years, 5 months ago

The correct answers are C. & D. Binary logging is a feature of MySQL that records all changes made to the database. This log can be used to restore the database to a previous state in case of a failure. Automated backups are regularly scheduled backups of the database. They are the most reliable way to ensure that data is not lost in case of a catastrophic failure.

upvoted 4 times

heretolearnazure 2 years, 5 months ago

C and D. Read replicas wont work in this case.

upvoted 1 times

FaizAhmed 2 years, 7 months ago

C. Binary logging D. Automated backups

upvoted 1 times

taer 2 years, 10 months ago

Binary logging records changes to the data, which can help you recover data and minimize data loss during an unexpected failure. Automated backups create regular backups of your database, allowing you to restore the database to a specific point in time in case of a catastrophic failure.

upvoted 3 times

abbottWang 2 years, 11 months ago

backup data automatically

upvoted 1 times

telp 2 years, 11 months ago

CD => the answer B is for performance issue. The question focus on data loss prevention.So you are going with a SINGLE instance of MySQL for a critical business application.

upvoted 1 times

telp 2 years, 11 months ago

CD => the answer B is for performance issue. The question focus on data loss prevention.

upvoted 1 times

medi01 2 years, 10 months ago

So you are going with a SINGLE instance of MySQL for a critical business application.

upvoted 1 times

okixavi 3 years ago

Selected Answer: BD

B and D: No need to explain D, but B... here is why When you set up a read replica, automaticaly binary logging is activated. Then, in case of desaster, you can promote manually a read replica and it will have all data before the desaster occurs.

upvoted 3 times

okixavi 3 years ago

B and D: No need to explain D, but B... here is why When you set up a read replica, automaticaly binary logging is activated. Then, in case of desaster, you can promote manually a read replica and it will have all data before the desaster occurs.sure, binary logging starts Automatically upon configuring read-replica?? - Don't think so, https://cloud.google.com/sql/docs/mysql/replication/create-replica

upvoted 3 times

r1ck 3 years ago

sure, binary logging starts Automatically upon configuring read-replica?? - Don't think so,
https://cloud.google.com/sql/docs/mysql/replication/create-replica

upvoted 2 times

Jeena345 3 years ago

B and D are correct answers as per below reference, 1. Before you can create a read replica of a primary Cloud SQL instance, the instance must meet the following requirements: Automated backups must be enabled. 2. Binary logging must be enabled which requires point-in-time recovery to be enabled. Learn more about the impact of these logs. 3. At least one backup must have been created after binary logging was enabled. It means creating read replica already covers binary logging. Please read the following references for more information https://cloud.google.com/solutions/cloud-sql-mysql-disaster-recovery-complete-failover-fallback https://medium.com/google-cloud/cloud-sql-recovering-from-regional-failure-in-10-minutes-or-less-mysql-fc055540a8f0 Replication in Cloud SQL | Cloud SQL for MySQL | Google CloudYes, you are correct that creating a read replica requires binary logging to be enabled on the primary instance. However, the purpose of a read replica is to scale read traffic and offload it from the primary instance, not to prevent data loss in case of catastrophic failure. While enabling binary logging is a requirement for creating a read replica, it is not the primary purpose of a read replica. IMO the two features that should be implemented to ensure minimum data loss in case of catastrophic failure are Binary logging and Automated backups.

upvoted 1 times

Jeena345 3 years ago

B and D are correct answers as per below reference, 1. Before you can create a read replica of a primary Cloud SQL instance, the instance must meet the following requirements:

Automated backups must be enabled.

2. Binary logging must be enabled which requires point-in-time recovery to be enabled. Learn more about the impact of these logs.

3. At least one backup must have been created after binary logging was enabled.

It means creating read replica already covers binary logging.

Please read the following references for more information

https://cloud.google.com/solutions/cloud-sql-mysql-disaster-recovery-complete-failover-fallback

https://medium.com/google-cloud/cloud-sql-recovering-from-regional-failure-in-10-minutes-or-less-mysql-fc055540a8f0

Replication in Cloud SQL | Cloud SQL for MySQL | Google Cloud

upvoted 1 times

mmathiou 2 years, 9 months ago

Yes, you are correct that creating a read replica requires binary logging to be enabled on the primary instance. However, the purpose of a read replica is to scale read traffic and offload it from the primary instance, not to prevent data loss in case of catastrophic failure. While enabling binary logging is a requirement for creating a read replica, it is not the primary purpose of a read replica. IMO the two features that should be implemented to ensure minimum data loss in case of catastrophic failure are Binary logging and Automated backups.

upvoted 3 times

examch 3 years, 1 month ago

C and D are correct answers, Backups help you restore lost data to your Cloud SQL instance. Additionally, if an instance is having a problem, you can restore it to a previous state by using the backup to overwrite it. Enable automated backups for any instance that contains necessary data. Backups protect your data from loss or damage. Enabling automated backups, along with binary logging, is also required for some operations, such as clone and replica creation. https://cloud.google.com/sql/docs/mysql/backup-recovery/backups#what_backups_provide

upvoted 2 times

omermahgoub 3 years, 1 month ago

C. Binary logging D. Automated backups Binary logging is a feature of MySQL that records all changes made to the database in a binary log file. By enabling binary logging on your Cloud SQL instance, you can use the log file to recover your database in case of catastrophic failure. Automated backups are a feature of Cloud SQL that allows you to automatically create and retain backups of your database. By enabling automated backups, you can restore your database in case of catastrophic failure or other data loss events. Option A, sharding, is not a recommended approach. Sharding is a technique for distributing data across multiple servers to improve performance and scalability. While sharding can help to improve the performance of a database, it is not specifically designed to protect against data loss in case of catastrophic failure.Option B, read replicas, is not a recommended approach. Read replicas are copies of a database that can be used to offload read traffic from the primary database. While read replicas can improve the performance of a database, they are not specifically designed to protect against data loss in case of catastrophic failure. Option E, semisynchronous replication, is not a recommended approach. Semisynchronous replication is a method of replicating data between a primary database and one or more secondary databases. While semisynchronous replication can help to ensure that data is replicated quickly and accurately, it is not specifically designed to protect against data loss in case of catastrophic failure.

upvoted 4 times

omermahgoub 3 years, 1 month ago

C. Binary logging D. Automated backups

Binary logging is a feature of MySQL that records all changes made to the database in a binary log file. By enabling binary logging on your Cloud SQL instance, you can use the log file to recover your database in case of catastrophic failure.

Automated backups are a feature of Cloud SQL that allows you to automatically create and retain backups of your database. By enabling automated backups, you can restore your database in case of catastrophic failure or other data loss events.

Option A, sharding, is not a recommended approach. Sharding is a technique for distributing data across multiple servers to improve performance and scalability. While sharding can help to improve the performance of a database, it is not specifically designed to protect against data loss in case of catastrophic failure.

upvoted 4 times

omermahgoub 3 years, 1 month ago

Option B, read replicas, is not a recommended approach. Read replicas are copies of a database that can be used to offload read traffic from the primary database. While read replicas can improve the performance of a database, they are not specifically designed to protect against data loss in case of catastrophic failure. Option E, semisynchronous replication, is not a recommended approach. Semisynchronous replication is a method of replicating data between a primary database and one or more secondary databases. While semisynchronous replication can help to ensure that data is replicated quickly and accurately, it is not specifically designed to protect against data loss in case of catastrophic failure.

upvoted 2 times

megumin 3 years, 3 months ago

cd is ok

upvoted 1 times

Question #109 *Topic 1*

You are working at a sports association whose members range in age from 8 to 30. The association collects a large amount of health data, such as sustained injuries. You are storing this data in BigQuery. Current legislation requires you to delete such information upon request of the subject. You want to design a solution that can accommodate such a request. What should you do?

A. Use a unique identifier for each individual. Upon a deletion request, delete all rows from BigQuery with this identifier.

B. When ingesting new data in BigQuery, run the data through the Data Loss Prevention (DLP) API to identify any personal information. As part of the DLP scan, save the result to Data Catalog. Upon a deletion request, query Data Catalog to find the column with personal information.

C. Create a BigQuery view over the table that contains all data. Upon a deletion request, exclude the rows that affect the subject's data from this view. Use this view instead of the source table for all analysis tasks.

D. Use a unique identifier for each individual. Upon a deletion request, overwrite the column with the unique identifier with a salted SHA256 of its value.

Suggested Answer: *B*

Official Answer: B

Community vote distribution



milan74 Highly Voted 5 years, 1 month ago

According to me, the question states "The association collects a large amount of health data, such as sustained injuries." and the nuance on the word such => " Current legislation requires you to delete "SUCH" information upon request of the subject. " So from that point of view the question is not to delete the entire user records but specific data related to personal health data. With DLP you can use InfoTypes and InfoType detectors to specifically scan for those entries and how to act upon them (link https://cloud.google.com/dlp/docs/concepts-infotypes) I would say B.There is no need of DLP. All the data is sensitive but only upon user request it needs deletion. So A should be the correct answer.as PhilipKoku mentioned below: A) is the correct answer. B) is only masking the data and then when a request is received, it identified the record but it doesn't delete it. D) Is masking the ID.B is not masking the data but identifying where it is to take action on at later date if requiredI want to delete all the informations about the user, not only those individuate by DLP. ALL THE INFORMATIONS of the users...B is not correct! the correct is A(A) - Primary task is "legislation requires you to delete" .. and B is not deleting. only A is deletingDeletion is implied in "Upon a deletion request, query Data Catalog to find the column with personal information."

upvoted 86 times

milan74 Highly Voted 5 years, 1 month ago

According to me, the question states "The association collects a large amount of health data, such as sustained injuries." and the nuance on the word such => " Current legislation requires you to delete "SUCH" information upon request of the subject. " So from that point of view the question is not to delete the entire user records but specific data related to personal health data. With DLP you can use InfoTypes and InfoType detectors to specifically scan for those entries and how to act upon them (link https://cloud.google.com/dlp/docs/concepts-infotypes) I would say B.

upvoted 86 times

lshu_awsguy 4 years ago

There is no need of DLP. All the data is sensitive but only upon user request it needs deletion. So A should be the correct answer.

upvoted 12 times

Arad 4 years, 8 months ago

as PhilipKoku mentioned below: A) is the correct answer. B) is only masking the data and then when a request is received, it identified the record but it doesn't delete it. D) Is masking the ID.B is not masking the data but identifying where it is to take action on at later date if required

upvoted 12 times

Arad 4 years, 8 months ago

as PhilipKoku mentioned below: A) is the correct answer. B) is only masking the data and then when a request is received, it identified the record but it doesn't delete it. D) Is masking the ID.

upvoted 12 times

mgm7 4 years, 8 months ago

B is not masking the data but identifying where it is to take action on at later date if required

upvoted 6 times

zanfo 4 years, 5 months ago

I want to delete all the informations about the user, not only those individuate by DLP. ALL THE INFORMATIONS of the users...B is not correct! the correct is A

upvoted 8 times

AmitAr 4 years, 3 months ago

(A) - Primary task is "legislation requires you to delete" .. and B is not deleting. only A is deleting

upvoted 12 times

AmitAr 4 years, 3 months ago

(A) - Primary task is "legislation requires you to delete" .. and B is not deleting. only A is deletingDeletion is implied in "Upon a deletion request, query Data Catalog to find the column with personal information."

upvoted 12 times

BeCalm 3 years, 5 months ago

Deletion is implied in "Upon a deletion request, query Data Catalog to find the column with personal information."

upvoted 2 times

XDevX 5 years, 1 month ago

IMHO a) is the correct answer because it is easier to operate. The question is not how to mask data and so on but just to delete data on request, so I don't think that we have to use for just the deletion of specific data DLP.

upvoted 37 times

Pime13 Most Recent 6 months ago

Direct Compliance: Using a unique identifier allows for a targeted DELETE statement in BigQuery that permanently removes specific records from the table, satisfying the legal requirement to "delete" the information. Data Integrity: Unique identifiers (like a user_id or individual_id) ensure that only the relevant data for the requesting subject is removed without impacting the data of other members. Operational Simplicity: Standard SQL DELETE with a WHERE clause is a well-supported DML (Data Manipulation Language) operation in BigQuery, making it the most straightforward design to implement for individual data removal.Option B (DLP API & Data Catalog): While Cloud Data Loss Prevention (DLP) is excellent for identifying or obfuscating PII/PHI during ingestion, simply saving the scan result to Data Catalog does not facilitate the actual physical deletion of existing data required by the prompt.

upvoted 1 times

Pime13 6 months ago

Selected Answer: A

Direct Compliance: Using a unique identifier allows for a targeted DELETE statement in BigQuery that permanently removes specific records from the table, satisfying the legal requirement to "delete" the information.

Data Integrity: Unique identifiers (like a user_id or individual_id) ensure that only the relevant data for the requesting subject is removed without impacting the data of other members.

Operational Simplicity: Standard SQL DELETE with a WHERE clause is a well-supported DML (Data Manipulation Language) operation in BigQuery, making it the most straightforward design to implement for individual data removal.

upvoted 1 times

Pime13 6 months ago

Option B (DLP API & Data Catalog): While Cloud Data Loss Prevention (DLP) is excellent for identifying or obfuscating PII/PHI during ingestion, simply saving the scan result to Data Catalog does not facilitate the actual physical deletion of existing data required by the prompt.

upvoted 1 times

CtrlX 11 months, 2 weeks ago

answer A. Was in my exam

upvoted 2 times

halifax 1 year, 5 months ago

Option A doesn't Delete all instances of PII, what happens if there is personal information on the descriptions? such as " sustained injuries". Delete just the identifiers is not complete.

upvoted 1 times

david_tay 1 year, 5 months ago

B is pointless as any health data is sensitive data, and also the fact that it did not mention deletion. Hence A is the correct answer.

upvoted 1 times

plumbig11 1 year, 7 months ago

Use a unique identifier for each individual. Upon a deletion request, delete all rows from BigQuery with this identifier.

upvoted 1 times

rrope 1 year, 7 months ago

A. Use a unique identifier for each individual.

upvoted 1 times

andyk87 1 year, 8 months ago

Option B is better when the requirement is to delete only the PII health data, not all data related to the individual.

upvoted 1 times

Sephethus 2 years, 2 months ago

It had better be A, if not then you're not a good organization

upvoted 2 times

hitmax87 2 years, 3 months ago

Data Loss Prevention must have!

upvoted 1 times

Gino17m 2 years, 3 months ago

I vote for B. I had some doubts whether A was correct, but: - I'm not convinced by the argument "only A talks about deleting" (it would be too easy if it was about choosing an answer containing the word "delete" ;) - the question says "design a solution that can accommodate such a request" - I'm not very fluent in english, but "accommodate" imho means more "facilitate" than "accomplish" here - I think that the task is about deleting health data not everything related with unique identifier - Data Catalog allows you to manage data, knowing in which datasets and in which tables what data is stored. Answer "A" somehow imposes the data model - each table with data related to a given individual must contain the ID of this individual (in a real data model this does not have to be the case).

upvoted 3 times

Djenko 2 years, 5 months ago

Should be A)

upvoted 2 times

mesodan 2 years, 5 months ago

A is correct. As for option B: While DLP is valuable for identifying sensitive data, it might not be sufficient for this specific case. DLP cannot necessarily determine an individual's right to deletion based solely on data classification. Additionally, relying on Data Catalog to store the results adds unnecessary complexity and potential inconsistencies.

upvoted 4 times

Gall 2 years, 6 months ago

B. The A removes all data, not SUCH only.

upvoted 2 times

NoCrapEva 2 years, 6 months ago

Ans. B assumes you will delete the Personal Information found in the Catalog... Some people are reading GDPR into this question (we are not told what country and what legislation). The question states you must delete all information (not just personal informarion) on request. Ans B is a red herring ! Answer must = A

upvoted 2 times

Roro_Brother 2 years, 7 months ago

(A) - Primary task is "legislation requires you to delete" .. and B is not deleting. only A is deleting

upvoted 5 times

Question #110 *Topic 1*

Your company has announced that they will be outsourcing operations functions. You want to allow developers to easily stage new versions of a cloud-based application in the production environment and allow the outsourced operations team to autonomously promote staged versions to production. You want to minimize the operational overhead of the solution. Which Google Cloud product should you migrate to?

A. App Engine

B. GKE On-Prem

C. Compute Engine

D. Google Kubernetes Engine

Suggested Answer: *D*

Reference: https://cloud.google.com/security/compliance/eba-outsourcing-mapping-gcp

Community vote distribution

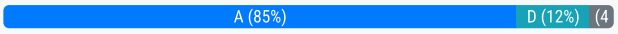

kopper2019 Highly Voted 5 years, 1 month ago

A. App Engine

upvoted 36 times

arsav Highly Voted 5 years ago

Answer should be A as only with App Engine we have a default service account which allows the user to deploy the changes per project. for GKE we may have to configure additional permission for both DEV and Operations team to deploy the changes.

https://cloud.google.com/appengine/docs/standard/php/service-account

upvoted 25 times

CtrlX Most Recent 11 months, 2 weeks ago

answer A. Was in my exam

upvoted 1 times

Sephethus 2 years, 2 months ago

A. App Engine. Explanation: Why A is correct: App Engine: Google App Engine is a fully managed platform-as-a-service (PaaS) that allows developers to build and deploy applications quickly and easily without worrying about managing the underlying infrastructure. It supports continuous integration and continuous deployment (CI/CD) processes, enabling developers to stage new versions of applications easily. Staging and Promotion: App Engine has built-in support for traffic splitting and versioning, which allows you to stage new versions of your application and gradually promote them to production. This can be done with minimal operational overhead, making it ideal for scenarios where operational functions are outsourced. Minimal Operational Overhead: Since App Engine is fully managed, it reduces the operational burden significantly, making it easier for the outsourced operations team to handle promotions and manage the application.

upvoted 2 times

JaimeMS 2 years, 2 months ago

A. App Engine

upvoted 1 times

jaisonPathiyil 2 years, 4 months ago

Is this answers are really correct or misleading to us..?

upvoted 3 times

mesodan 2 years, 5 months ago

While both GKE and App Engine offer functionalities for deploying cloud-based applications, App Engine is more managed service compared to GKE, resulting in lower operational overhead.

upvoted 2 times

Anandmrk 2 years, 5 months ago

I did my exam today and saw this question. But I am sure A was the answer due to the operational overhead phrase

upvoted 3 times

thewalker 2 years, 9 months ago

In the question, we see "cloud-based application" - I assume, it means cloud native -> dockers/containers -> K8s -> GKE. Hence, D is my option.

upvoted 2 times

AdityaGupta 2 years, 10 months ago

I agreed with "omermahgoub" the answer should be D. As you will bundle the application and its dependencies into container image and deploy. All environments will have same image deployed from Dev, TEST, Staging to PROD. There will be less operational overheard for operations team.

upvoted 1 times

nocrush 2 years, 11 months ago

A App Engine reduces ops overhead

upvoted 2 times

Frusci 2 years, 11 months ago

A. You deploy your new version to App Engine without setting it as the default version. The ops team then just has to make it the default version when they want to promote it. Simplest answer.

upvoted 2 times

heretolearnazure 2 years, 11 months ago

App Engine because of less ovehead.

upvoted 1 times

JohnWick2020 3 years, 2 months ago

Answer is A. By process of elimination you arrive at App Engine or GKE. Now the requirement is to "to minimize the operational overhead of the solution". On the IaaS to PaaS spectrum, this can only be App Engine! IaaS = Compute Engine. Hybrid = GKE (engineering heavy). PaaS = App Engine.

upvoted 3 times

Atanu 3 years, 2 months ago

"You want to minimize the operational overhead of the solution" ..This sentence is they key to go with Option A. GKE carries overhead as it's not purely PaaS.

upvoted 3 times

Sur_Nikki 3 years, 3 months ago

It should be D. As GKE is considered to be the master product/service for creating a deployment and managing and keeping all the environments in SYNC

upvoted 1 times

sithin_nair 3 years, 5 months ago

A is right as the requirement is to deploy new changes and manage the application with no operational overhead.

upvoted 1 times

Question #111 *Topic 1*

Your company is running its application workloads on Compute Engine. The applications have been deployed in production, acceptance, and development environments. The production environment is business-critical and is used 24/7, while the acceptance and development environments are only critical during office hours. Your CFO has asked you to optimize these environments to achieve cost savings during idle times. What should you do?

A. Create a shell script that uses the gcloud command to change the machine type of the development and acceptance instances to a smaller machine type outside of office hours. Schedule the shell script on one of the production instances to automate the task.

B. Use Cloud Scheduler to trigger a Cloud Function that will stop the development and acceptance environments after office hours and start them just before office hours.

C. Deploy the development and acceptance applications on a managed instance group and enable autoscaling.

D. Use regular Compute Engine instances for the production environment, and use preemptible VMs for the acceptance and development environments.

Suggested Answer: *B*

Reference: https://cloud.google.com/blog/products/it-ops/best-practices-for-optimizing-your-cloud-costs

Community vote distribution

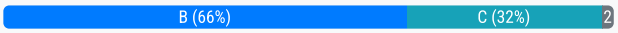

pamepadero Highly Voted 5 years, 1 month ago

B is the answer.

https://cloud.google.com/blog/products/it-ops/best-practices-for-optimizing-your-cloud-costs Schedule VMs to auto start and stop: The benefit of a platform like Compute Engine is that you only pay for the compute resources that you use. Production systems tend to run 24/7; however, VMs in development, test or personal environments tend to only be used during business hours, and turning them off can save you a lot of money!

https://cloud.google.com/blog/products/storage-data-transfer/save-money-by-stopping-and-starting-compute-engine-instances-on-schedule

Cloud Scheduler, GCP's fully managed cron job scheduler, provides a straightforward solution for automatically stopping and starting VMs. By employing Cloud Scheduler with Cloud Pub/Sub to trigger Cloud Functions on schedule, you can stop and start groups of VMs identified with labels of your choice (created in Compute Engine). Here you can see an example schedule that stops all VMs labeled "dev" at 5pm and restarts them at 9am, while leaving VMs labeled "prod" untouched

upvoted 40 times

pamepadero Highly Voted 5 years, 1 month ago

B is the answer. https://cloud.google.com/blog/products/it-ops/best-practices-for-optimizing-your-cloud-costs Schedule VMs to auto start and stop: The benefit of a platform like Compute Engine is that you only pay for the compute resources that you use. Production systems tend to run 24/7; however, VMs in development, test or personal environments tend to only be used during business hours, and turning them off can save you a lot of money! https://cloud.google.com/blog/products/storage-data-transfer/save-money-by-stopping-and-starting-compute-engine-instances-on-schedule Cloud Scheduler, GCP's fully managed cron job scheduler, provides a straightforward solution for automatically stopping and starting VMs. By employing Cloud Scheduler with Cloud Pub/Sub to trigger Cloud Functions on schedule, you can stop and start groups of VMs identified with labels of your choice (created in Compute Engine). Here you can see an example schedule that stops all VMs labeled "dev" at 5pm and restarts them at 9am, while leaving VMs labeled "prod" untouchedQuestion says that dev/test are "not critical", it doesn't mean that they are not needed at all ...Great answer and documentation. Def BExcellent .....even the good CFO is telling leave the office after 5.oo and come next day to work :)

upvoted 40 times

rzygor 4 years ago

Question says that dev/test are "not critical", it doesn't mean that they are not needed at all ...

upvoted 20 times

Ric350 4 years, 1 month ago

Great answer and documentation. Def B

upvoted 2 times

sgofficial 4 years ago

Excellent .....even the good CFO is telling leave the office after 5.oo and come next day to work :)

upvoted 16 times

kopper2019 5 years, 1 month ago

Ans) B , assuming VM doesn't need to be up after office hours .

upvoted 26 times

plumbig11 1 year, 7 months ago

Cloud scheduler considering the requisites.

upvoted 1 times

rrope 1 year, 7 months ago

Selected Answer: D

By utilizing preemptible VMs for the acceptance and development environments, you can significantly reduce costs without compromising the availability and performance of your critical production environment.

upvoted 1 times

rrope 1 year, 7 months ago

By utilizing preemptible VMs for the acceptance and development environments, you can significantly reduce costs without compromising the availability and performance of your critical production environment.problem is you won't know what time preemptible VMs will shut down. There is only a brief warning from Google. Hence not suitable since they need to be up during work hours.

upvoted 1 times

david_tay Most Recent 1 year, 5 months ago

problem is you won't know what time preemptible VMs will shut down. There is only a brief warning from Google. Hence not suitable since they need to be up during work hours.

upvoted 1 times

25lion52 1 year, 10 months ago

Stop the dev and acceptance envs is super weird. Any critical problems or overtimes will be an issue with this approach. Simple auto scaling environment is a good solution IMHO

upvoted 3 times

Gino17m 2 years, 4 months ago

B is right answer

upvoted 1 times

dija123 2 years, 4 months ago

Agree with B

upvoted 1 times

spuyol 2 years, 6 months ago

Answer D A: too complex and maybe small or zero saving if you can't find a valid smaller machine type B: Not valid. Question says that PRE environments are not critical after office hours. But it doesn't say no service at all C: Some risk is introduced if you have different architecture on PRE than PRO envs D: It's the only valid and realiable option. Simple and effective. It's my choice. In a real scenario I will first start with this and then review if the savings are enough before more complicated choicesAd. "B: Not valid. Question says that PRE environments are not critical after office hours. But it doesn't say no service at all" But the Question says that PRE environments are critical during office hours, so you can't use preemptible VMs - "Compute Engine might stop (preempt) these instances if it needs to reclaim the compute capacity for allocation to other VMs"

upvoted 3 times

spuyol 2 years, 6 months ago

Answer D

A: too complex and maybe small or zero saving if you can't find a valid smaller machine type B: Not valid. Question says that PRE environments are not critical after office hours. But it doesn't say no service at all C: Some risk is introduced if you have different architecture on PRE than PRO envs D: It's the only valid and realiable option. Simple and effective. It's my choice. In a real scenario I will first start with this and then review if the savings are enough before more complicated choices

upvoted 3 times

Gino17m 2 years, 4 months ago

Ad. "B: Not valid. Question says that PRE environments are not critical after office hours. But it doesn't say no service at all" But the Question says that PRE environments are critical during office hours, so you can't use preemptible VMs - "Compute Engine might stop (preempt) these instances if it needs to reclaim the compute capacity for allocation to other VMs"

upvoted 1 times

the1dv 2 years, 7 months ago

Selected Answer: C

MIG's with autoscaling will scale to Zero if not needed

upvoted 4 times

the1dv 2 years, 7 months ago

MIG's with autoscaling will scale to Zero if not neededsome risks are added if you have different architecture on PRO and PRE envs

upvoted 4 times

spuyol 2 years, 6 months ago

some risks are added if you have different architecture on PRO and PRE envs

upvoted 1 times

AWS_Sam 2 years, 7 months ago

B for sure

upvoted 1 times

parthkulkarni998 2 years, 8 months ago

Also managed instance group reduces instances in case of low/no-traffic incurring lesser charges. Ideally, its a cleaner approach considering the ask is to optimize during "idle time". Incase people are working in different time zones, late shifts it doesnt make sense to trigger shutdown at a predefined times.

upvoted 4 times

odacir 2 years, 9 months ago

B is the answer. But today, you don't need complicated CRON + CF. Auto shutdown by cron expression it's a feature built in: https://cloud.google.com/compute/docs/instances/schedule-instance-start-stop

upvoted 2 times

werdy92 2 years, 10 months ago

really wondering why not C...Not critical is not equivalent with not running at all....

upvoted 5 times

werdy92 2 years, 10 months ago

really wondering why not C...Not critical is not equivalent with not running at all....Also managed instance group reduces instances in case of low/no-traffic incurring lesser charges. Ideally, its a cleaner approach considering the ask is to optimize during "idle time". Incase people are working in different time zones, late shifts it doesnt make sense to trigger shutdown at a predefined times.

upvoted 5 times

wooloo 3 years ago

"are only critical during office hours" does not mean it could be completely stopped. So may the option C correct?

upvoted 6 times

mifrah 3 years, 4 months ago

In my opinion B is over-engineered: Why not just add an "instance schedule" for start/stop the Compute Engines? Why creating a scheduler and writing a Cloud Function...Just exactly what I have thought. It is enough with instance schedule". But GCP wants you to spend money and use cloud functions LOL

upvoted 3 times

mifrah 3 years, 4 months ago

In my opinion B is over-engineered: Why not just add an "instance schedule" for start/stop the Compute Engines? Why creating a scheduler and writing a Cloud Function...

upvoted 3 times

ccpmad 2 years, 2 months ago

Just exactly what I have thought. It is enough with instance schedule". But GCP wants you to spend money and use cloud functions LOL

upvoted 1 times

MaryMei 3 years, 5 months ago

https://cloud.google.com/compute/docs/instances/viewing-and-applying-idle-vm-recommendations B seems close to this Google provided service option, the extra step should be using idle VM recommendations to find and stop idle VM instances to reduce waste of resources

upvoted 1 times

PAUGURU 3 years, 6 months ago

Since the price of preemptibles is 1/4 the price of a standard machine D costs far less than B since office hours are 1/3 of whole day. It costs less to keep them running 24h as preemptibles.Yes, but preemptibles use GCP excess resources so you will achieve the opposite of the desired effect, during office hours, they will underperform in the best case (worst case will stop altogether) and, during non-office hours, preemptibles will work well...

upvoted 3 times

PAUGURU 3 years, 6 months ago

Since the price of preemptibles is 1/4 the price of a standard machine D costs far less than B since office hours are 1/3 of whole day. It costs less to keep them running 24h as preemptibles.

upvoted 3 times

DevOpsifier 3 years, 2 months ago

Yes, but preemptibles use GCP excess resources so you will achieve the opposite of the desired effect, during office hours, they will underperform in the best case (worst case will stop altogether) and, during non-office hours, preemptibles will work well...

upvoted 1 times

Question #112 *Topic 1*

You are moving an application that uses MySQL from on-premises to Google Cloud. The application will run on Compute Engine and will use Cloud SQL. You want to cut over to the Compute Engine deployment of the application with minimal downtime and no data loss to your customers. You want to migrate the application with minimal modification. You also need to determine the cutover strategy. What should you do?

A. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server. 2. Stop the on-premises application. 3. Create a mysqldump of the on-premises MySQL server. 4. Upload the dump to a Cloud Storage bucket. 5. Import the dump into Cloud SQL. 6. Modify the source code of the application to write queries to both databases and read from its local database. 7. Start the Compute Engine application. 8. Stop the on-premises application.

B. 1. Set up Cloud SQL proxy and MySQL proxy. 2. Create a mysqldump of the on-premises MySQL server. 3. Upload the dump to a Cloud Storage bucket. 4. Import the dump into Cloud SQL. 5. Stop the on-premises application. 6. Start the Compute Engine application.

C. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server. 2. Stop the on-premises application. 3. Start the Compute Engine application, configured to read and write to the on-premises MySQL server. 4. Create the replication configuration in Cloud SQL. 5. Configure the source database server to accept connections from the Cloud SQL replica. 6. Finalize the Cloud SQL replica configuration. 7. When replication has been completed, stop the Compute Engine application. 8. Promote the Cloud SQL replica to a standalone instance. 9. Restart the Compute Engine application, configured to read and write to the Cloud SQL standalone instance.

D. 1. Stop the on-premises application. 2. Create a mysqldump of the on-premises MySQL server. 3. Upload the dump to a Cloud Storage bucket. 4. Import the dump into Cloud SQL. 5. Start the application on Compute Engine.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
C (100%)

victory108 Highly Voted 4 years, 7 months ago
C. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server. 2. Stop the on-premises application. 3. Start the Compute Engine application, configured to read and write to the on-premises MySQL server. 4. Create the replication configuration in Cloud SQL. 5. Configure the source database server to accept connections from the Cloud SQL replica. 6. Finalize the Cloud SQL replica configuration. 7. When replication has been completed, stop the Compute Engine application. 8. Promote the Cloud SQL replica to a standalone instance. 9. Restart the Compute Engine application, configured to read and write to the Cloud SQL standalone instance.
upvoted 33 times

victory108 Highly Voted 4 years, 7 months ago
C. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server. 2. Stop the on-premises application. 3. Start the Compute Engine application, configured to read and write to the on-premises MySQL server. 4. Create the replication configuration in Cloud SQL. 5. Configure the source database server to accept connections from the Cloud SQL replica. 6. Finalize the Cloud SQL replica configuration. 7. When replication has been completed, stop the Compute Engine application. 8. Promote the Cloud SQL replica to a standalone instance. 9. Restart the Compute Engine application, configured to read and write to the Cloud SQL standalone instance.Agree with C. The only confusing is step "5. Configure the source database server to accept connections from the Cloud SQL replica." Is that not replication should go in the opposite direction from the on-premise (a.k.a. "source") database to Cloud SQL replica (presuming the latter is configured with a public IP address)?
upvoted 33 times

don_v 2 years, 1 month ago
Agree with C. The only confusing is step "5. Configure the source database server to accept connections from the Cloud SQL replica." Is that not replication should go in the opposite direction from the on-premise (a.k.a. "source") database to Cloud SQL replica (presuming the latter is configured with a public IP address)?
upvoted 1 times

kopper2019 4 years, 7 months ago
Ans C, from this guy muhasinem External replica promotion migration In the migration strategy of external replica promotion, you create an external database replica and synchronize the existing data to that replica. This can happen with minimal downtime to the existing database. When you have a replica database, the two databases have different roles that are referred to in this document as primary and replica. After the data is synchronized,

you promote the replica to be the primary in order to move the management layer with minimal impact to database uptime. In Cloud SQL, an easy way to accomplish the external replica promotion is to use the automated migration workflow. This process automates many of the steps that are needed for this type of migration.

upvoted 23 times

de1001c Most Recent 1 year, 8 months ago

Minimal downtime. Downtime in A is time to take mysql dump + fix potential failures, not good. Downtime in C is just the time from restart the service.

upvoted 1 times

Gino17m 1 year, 10 months ago

I wonder how Examtopics determines the so-called "Correct Answer" ..... C is correct

upvoted 1 times

heretolearnazure 2 years, 5 months ago

Answer is C

upvoted 1 times

aliounegdiop 2 years, 8 months ago

B is the correct answ

upvoted 2 times

BiddlyBdoyng 2 years, 10 months ago

Option A, writing to two databases form the app :(Option C all the way, it also aligns to GCP Data Migration Service.

upvoted 2 times

DRK8109 2 years, 10 months ago

mysql dump always causes long downtime.

upvoted 4 times

musumusu 2 years, 10 months ago

Correct Answer A C is unnecceory expensive and loss of data at after step 3

upvoted 2 times

BeCalm 2 years, 11 months ago

C seems to be the best answer but it is still a bit confusing. So basically there's a bi-directional sync between the 2 databases? Cloud instance is the primary and is writing into the on-prem and on-prem is being replicated into the Cloud.

upvoted 2 times

NodummyIQ 3 years, 1 month ago

Option C is not the correct answer because it involves modifying the application to read and write to both the on-premises MySQL server and Cloud SQL, which would involve significant modification to the application and could introduce potential complications or errors. It is generally better to minimize modification to the application when performing a migration. Option D, on the other hand, involves simply importing a mysqldump of the on-premises MySQL server into Cloud SQL and starting the application on Compute Engine, which is a simpler and more straightforward approach that involves minimal modification to the application.

upvoted 2 times

SureshbabuK 3 years, 1 month ago

Examtopic providing A as correct answer is causing confusion,

upvoted 6 times

Jose56 3 years, 2 months ago

C for minimal downtime

upvoted 2 times

megumin 3 years, 3 months ago

C is ok

upvoted 1 times

zr79 3 years, 4 months ago

Answer is C we have a new service https://cloud.google.com/database-migration

upvoted 7 times

minmin2020 3 years, 4 months ago

C because it has minimal modification to the application or database. Also it's easier to fail back to the original solution if the cloud implementation has issues (assuming that there will be a "post-go-live" monitoring period).

upvoted 2 times

Question #113 *Topic 1*

Your organization has decided to restrict the use of external IP addresses on instances to only approved instances. You want to enforce this requirement across all of your Virtual Private Clouds (VPCs). What should you do?

A. Remove the default route on all VPCs. Move all approved instances into a new subnet that has a default route to an internet gateway.

B. Create a new VPC in custom mode. Create a new subnet for the approved instances, and set a default route to the internet gateway on this new subnet.

C. Implement a Cloud NAT solution to remove the need for external IP addresses entirely.

D. Set an Organization Policy with a constraint on constraints/compute.vmExternalIpAccess. List the approved instances in the allowedValues list.

Suggested Answer: *D*

Reference: https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address

Community vote distribution



victory108 Highly Voted 4 years, 7 months ago

D. Set an Organization Policy with a constraint on constraints/compute.vmExternalIpAccess. List the approved instances in the allowedValues list.

upvoted 24 times

AnilKr Highly Voted 4 years, 6 months ago

Ans - D, https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#disableexternalip you might want to restrict external IP address so that only specific VM instances can use them. This option can help to prevent data exfiltration or maintain network isolation. Using an Organization Policy, you can restrict external IP addresses to specific VM instances with constraints to control use of external IP addresses for your VM instances within an organization or a project.

upvoted 21 times

CtrlX Most Recent 11 months, 2 weeks ago

answer D. Was in my exam

upvoted 1 times

plumbig11 1 year, 7 months ago

to enforce this requirement across all of both Virtual Private Clouds (VPCs) you have to set an organization policy.

upvoted 1 times

james2033 1 year, 8 months ago

https://cloud.google.com/compute/docs/ip-addresses/configure-static-external-ip-address#disableexternalip

upvoted 1 times

odacir 1 year, 11 months ago

"You cannot apply the constraint retroactively. All VMs that have external IP addresses before you enable the policy retain their external IP addresses." https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#disableexternalip It shouldn't be option D then

upvoted 1 times

beehive 3 years, 1 month ago

D is correct one. https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#disableexternalip

upvoted 3 times

rascalbrick 3 years, 2 months ago

Show on my Exam,unfortunatekt im failed..:(

upvoted 2 times

megumin 3 years, 3 months ago

D is ok

upvoted 1 times

AzureDP900 3 years, 4 months ago

D is correct

upvoted 1 times

2M 3 years, 5 months ago

option D

upvoted 2 times

AzureDP900 3 years, 7 months ago

D is right. constraints/compute.vmExternalIpAccess

upvoted 1 times

JoeyCASD 3 years, 9 months ago

vote for D https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#disableexternalip

upvoted 2 times

ss909098 3 years, 11 months ago

D it is

upvoted 1 times

technodev 4 years, 1 month ago

Got this question in my exam, answered D

upvoted 2 times

haroldbenites 4 years, 2 months ago

Go for D. https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#disableexternalip

upvoted 2 times

vincy2202 4 years, 2 months ago

D is the correct answer https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#disableexternalip

upvoted 1 times

Question #114 *Topic 1*

Your company uses the Firewall Insights feature in the Google Network Intelligence Center. You have several firewall rules applied to Compute Engine instances.
You need to evaluate the efficiency of the applied firewall ruleset. When you bring up the Firewall Insights page in the Google Cloud Console, you notice that there are no log rows to display. What should you do to troubleshoot the issue?

A. Enable Virtual Private Cloud (VPC) flow logging.

B. Enable Firewall Rules Logging for the firewall rules you want to monitor.

C. Verify that your user account is assigned the compute.networkAdmin Identity and Access Management (IAM) role.

D. Install the Google Cloud SDK, and verify that there are no Firewall logs in the command line output.

Suggested Answer: *B*
Reference: https://cloud.google.com/network-intelligence-center/docs/firewall-insights/how-to/using-firewall-insights

Community vote distribution
B (100%)

nohel Highly Voted 5 years, 1 month ago
Answer is B when you create a firewall rule there is an option for firewall rule logging on/off. It is set to off by default. To get firewall insights or view the logs for a specific firewall rule you need to enable logging while creating the rule or you can enable it by editing that rule.
https://cloud.google.com/network-intelligence-center/docs/firewall-insights/how-to/using-firewall-insights#enabling-fw-rules-logging
upvoted 37 times

victory108 Highly Voted 5 years, 1 month ago
B. Enable Firewall Rules Logging for the firewall rules you want to monitor.
upvoted 15 times

plumbig11 Most Recent 1 year, 7 months ago
Insights of firewall rules, Firewall rules logging is the best pratice, once you have a good logging and monitoring tool. B. Enable Firewall Rules Logging for the firewall rules you want to monitor.
upvoted 1 times

GlebG 2 years, 1 month ago
First D, then B
upvoted 1 times

Gino17m 2 years, 4 months ago
Corrent answer is B
upvoted 2 times

RVivek 3 years, 6 months ago
https://cloud.google.com/vpc/docs/firewall-rules-logging
upvoted 1 times

windsor_43 3 years, 7 months ago
The Answer is B Just had my exam today with a pass, this question was in the exam. Dated 31/12/22 Thanks to this site it was by far my most valuable
upvoted 5 times

jay9114 3 years, 8 months ago
You have to enable logging for a firewall rule in order to see the rows. "When you enable logging for a firewall rule, Google Cloud creates an entry called a connection record each time the rule allows or denies traffic." https://cloud.google.com/vpc/docs/firewall-rules-logging
upvoted 1 times

megumin 3 years, 9 months ago
B is ok
upvoted 1 times

minmin2020 3 years, 10 months ago
B. Enable Firewall Rules Logging for the firewall rules you want to monitor.
upvoted 1 times

AzureDP900 3 years, 10 months ago

Enable firewall rules logging , B is right

upvoted 1 times

DrishaS4 4 years ago

https://cloud.google.com/network-intelligence-center/docs/firewall-insights/how-to/using-firewall-insights#enabling-fw-rules-logging

upvoted 2 times

AzureDP900 4 years, 1 month ago

B is most appropriate answer, I will choose B.https://cloud.google.com/vpc/docs/firewall-rules-logging

upvoted 1 times

AzureDP900 4 years, 1 month ago

B is most appropriate answer, I will choose B.

upvoted 1 times

AzureDP900 4 years, 1 month ago

https://cloud.google.com/vpc/docs/firewall-rules-logging

upvoted 1 times

tannV 4 years, 3 months ago

Answered B. Got this question!

upvoted 2 times

azureaspirant 4 years, 6 months ago

02/15/21 exam

upvoted 1 times

haroldbenites 4 years, 8 months ago

Go for B

upvoted 1 times

vincy2202 4 years, 8 months ago

B is the correct answer https://cloud.google.com/network-intelligence-center/docs/firewall-insights/how-to/using-firewall-insights

upvoted 1 times

Question #115 *Topic 1*

Your company has sensitive data in Cloud Storage buckets. Data analysts have Identity Access Management (IAM) permissions to read the buckets. You want to prevent data analysts from retrieving the data in the buckets from outside the office network. What should you do?

A. 1. Create a VPC Service Controls perimeter that includes the projects with the buckets. 2. Create an access level with the CIDR of the office network.

B. 1. Create a firewall rule for all instances in the Virtual Private Cloud (VPC) network for source range. 2. Use the Classless Inter-domain Routing (CIDR) of the office network.

C. 1. Create a Cloud Function to remove IAM permissions from the buckets, and another Cloud Function to add IAM permissions to the buckets. 2. Schedule the Cloud Functions with Cloud Scheduler to add permissions at the start of business and remove permissions at the end of business.

D. 1. Create a Cloud VPN to the office network. 2. Configure Private Google Access for on-premises hosts.

Suggested Answer: *C*
Official Answer: C

Community vote distribution
A (100%)

TotoroChina Highly Voted 4 years, 7 months ago
Should be A. For all Google Cloud services secured with VPC Service Controls, you can ensure that: Resources within a perimeter are accessed only from clients within authorized VPC networks using Private Google Access with either Google Cloud or on-premises. https://cloud.google.com/vpc-service-controls/docs/overviewCorrect, this is about data exfiltration. See: https://youtu.be/EXwJFL24QzYThanks for including the youtube video it was very helpfulnice one, thank you manEnforce a security perimeter with VPC Service Controls to isolate resources of multi-tenant Google Cloud services—reducing the risk of data exfiltration or data breach.
upvoted 73 times

TotoroChina Highly Voted 4 years, 7 months ago
Should be A. For all Google Cloud services secured with VPC Service Controls, you can ensure that: Resources within a perimeter are accessed only from clients within authorized VPC networks using Private Google Access with either Google Cloud or on-premises. https://cloud.google.com/vpc-service-controls/docs/overview
upvoted 73 times

poseidon24 4 years, 6 months ago
Correct, this is about data exfiltration. See: https://youtu.be/EXwJFL24QzY
upvoted 14 times

poseidon24 4 years, 6 months ago
Correct, this is about data exfiltration. See: https://youtu.be/EXwJFL24QzYThanks for including the youtube video it was very helpfulnice one, thank you man
upvoted 14 times

mv2000 3 years, 7 months ago
Thanks for including the youtube video it was very helpful
upvoted 1 times

Sivanaga 3 years, 5 months ago
nice one, thank you man
upvoted 2 times

ArtistS 2 years, 3 months ago
Enforce a security perimeter with VPC Service Controls to isolate resources of multi-tenant Google Cloud services—reducing the risk of data exfiltration or data breach.
upvoted 1 times

XDevX 4 years, 7 months ago
IMHO c is wrong - the question is not to restrict access only for business hours but to restrict access to office network. In my opinion the only realistic approach seems to be a) https://cloud.google.com/vpc-service-controls/docs/supported-products#table_storage
upvoted 17 times

CtrlX Most Recent 11 months, 2 weeks ago

answer A. Was in my exam

upvoted 1 times

Xpzebra 1 year, 4 months ago

A. 1. [illegible] VPC [illegible]2. [illegible] CIDR □□□□□□□

upvoted 1 times

plumbig11 1 year, 7 months ago

In this case a VPC Service Controls , is more appropriated.

upvoted 1 times

19040e5 1 year, 9 months ago

It's obviously A. C mentions office hours which has nothing to do with the question!

upvoted 1 times

Gino17m 1 year, 10 months ago

A is correct answer. Examtopics should change so-called "Correct Answer" from C to A to stop confusin users.

upvoted 2 times

kalyan_krishna742020 1 year, 10 months ago

I'm preparing for a test and see that questions from 115 onwards are considered valid. Can anyone who's taken the test offer any insights or advice? Thank you!

upvoted 1 times

discuss24 2 years, 1 month ago

A is the correct answer. The question is specific to accessing the data outside of the office network. If the question talked about outside of work business hours then, we can consider C

upvoted 1 times

JPA210 2 years, 4 months ago

How can be possible that Examtopics say that the correct answer is C?! It doesn't make any sense! A is the correct one.

upvoted 2 times

heretolearnazure 2 years, 5 months ago

A is correct answer

upvoted 1 times

RVivek 3 years ago

https://cloud.google.com/vpc-service-controls/docs/overview

upvoted 1 times

vamgcp 3 years ago

A is correct because, For all Google Cloud services secured with VPC Service Controls, you can ensure that resources within a perimeter are accessed only from clients within authorized VPC networks using Private Google Access with either Google Cloud or on-premises.

upvoted 1 times

examch 3 years, 1 month ago

A is the correct answer, https://cloud.google.com/vpc-service-controls/docs/overview#isolate * A VM within a Virtual Private Cloud (VPC) network that is part of a service perimeter can read from or write to a Cloud Storage bucket in the same perimeter. However, VPC Service Controls doesn't allow VMs within VPC networks that are outside the perimeter to access Cloud Storage buckets that are inside the perimeter. * A copy operation between two Cloud Storage buckets succeeds if both buckets are in the same service perimeter, but if one of the buckets is outside the perimeter, the copy operation fails. * VPC Service Controls doesn't allow a VM within a VPC network that is inside a service perimeter to access Cloud Storage buckets that are outside the perimeter.

upvoted 4 times

thamaster 3 years, 1 month ago

answer C will not prevent connection from outside of office network

upvoted 1 times

cshubham173 3 years, 2 months ago

For all Google Cloud services secured with VPC Service Controls, you can ensure that: Resources within a perimeter are accessed only from clients within authorized VPC networks using Private Google Access with either Google Cloud or on-premises. https://cloud.google.com/vpc-service-controls/docs/overview

upvoted 2 times

megumin 3 years, 3 months ago

A is ok

upvoted 1 times

Question #116 *Topic 1*

You have developed a non-critical update to your application that is running in a managed instance group, and have created a new instance template with the update that you want to release. To prevent any possible impact to the application, you don't want to update any running instances. You want any new instances that are created by the managed instance group to contain the new update. What should you do?

A. Start a new rolling restart operation.

B. Start a new rolling replace operation.

C. Start a new rolling update. Select the Proactive update mode.

D. Start a new rolling update. Select the Opportunistic update mode.

Suggested Answer: *C*

Official Answer: C

Community vote distribution

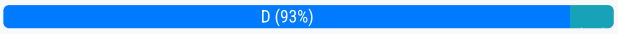

XDevX Highly Voted 5 years, 1 month ago

IMHO the correct answer is d) opportunistic mode, not c) proactive mode. The requirement is not to update any running instances. see: https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups For automated rolling updates, you must set the mode to proactive. Alternatively, if an automated update is potentially too disruptive, you can choose to perform an opportunistic update. The MIG applies an opportunistic update only when you manually initiate the update on selected instances or when new instances are created. New instances can be created when you or another service, such as an autoscaler, resizes the MIG.

upvoted 60 times

victory108 Highly Voted 5 years, 1 month ago

D. Start a new rolling update. Select the Opportunistic update mode.

upvoted 12 times

Chang401 Most Recent 1 year, 11 months ago

In Google Cloud, the main difference between proactive and opportunistic updates in a managed instance group (MIG) is when they are applied: Proactive updates: Automatically apply updates to existing VMs. Opportunistic updates: Only apply updates when you manually select a VM to update or when new instances are created. for all those who thinks its C please google and check its a straight forward question and you guyz are confusing people

upvoted 1 times

shashii82 2 years, 5 months ago

Option C. To release a non-critical update to your application without updating any running instances and ensuring that new instances created by the managed instance group contain the new update, you should choose option C: Start a new rolling update and select the Proactive update mode. In the Proactive update mode, the managed instance group creates new instances with the updated template while keeping the existing instances running until they are eventually replaced. This allows you to roll out the update gradually without affecting the currently running instances.

upvoted 4 times

kip21 2 years, 7 months ago

D - Correct Managed instance groups support two types of update: Automatic, or proactive, updates Selective, or opportunistic, updates If you want to apply updates automatically, set the type to proactive. Alternatively, if an automated update is potentially too disruptive, you can choose to perform an opportunistic update. The MIG applies an opportunistic update only when you manually initiate the update on selected instances or when new instances are created.

upvoted 3 times

discuss24 2 years, 7 months ago

D is correct, per Google's documentation (The MIG applies an opportunistic update only when you manually initiate the update on selected instances or when new instances are created. New instances can be created when you or another service, such as an autoscaler, resizes the MIG.)

upvoted 1 times

spuyol 2 years, 8 months ago

Selected answer: D "To prevent any possible impact to the application, you don't want to update any running instances" this is automatic achievable only by using opportunistic applying it during autoscale actions as documentation says: if you want to selectively apply a new configuration only to new or to specific instances in a MIG, see Selectively apply VM configuration updates in a MIG

upvoted 2 times

6b13108 2 years, 8 months ago

In my opinion C is correct for Proactive update mode. Considering the following doc; "....Automated updates support up to two instance template versions in your MIG. This means that you can specify two different instance template versions for your group, which is useful for performing canary updates. To start a basic rolling update where the update is applied to all instances in the group, follow the instructions below....." See Also: "Update type Managed instance groups support two types of update: Automatic, or proactive, updates Selective, or opportunistic, updates If you want to apply updates automatically, set the type to proactive." https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups#starting_a_basic_rolling_update

upvoted 2 times

thewalker 2 years, 9 months ago

https://cloud.google.com/compute/docs/instance-groups/updating-migs#choosing_between_automated_and_selective_updates Automatic (proactive): Use this method if you want the MIG to automatically apply new configurations to all or to a subset of existing VMs in the group. The level of disruption to running VMs depends on the update policy that you configure. You can use this method to canary update new instance templates. To use this method, set the MIG's update type to "proactive". Selective (opportunistic): Use this method if you want to apply the update manually or if you want to update all existing VMs in the group at once. You target any or all VMs to be updated to the latest configuration. To use this method, set the MIG's update type to "opportunistic". Hence, C

upvoted 3 times

JPA210 2 years, 10 months ago

It is option C, with proactive update, you are not updating the running instances, you start new ones with the new configuration template. And stop the old ones, so there is not disruption to the service.

upvoted 2 times

AdityaGupta 2 years, 10 months ago

To apply an updated configuration to existing VMs, you can set up an automatic update–also known as a proactive update type. The MIG automatically rolls out configuration updates to all or to a subset of the group's VMs. Alternatively, if an automated update is potentially too disruptive, you can choose to perform an opportunistic update. The MIG applies an opportunistic update only when you manually initiate the update on selected instances or when new instances are created. New instances can be created when you or another service, such as an autoscaler, resizes the MIG. https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups

upvoted 2 times

duzapo 2 years, 11 months ago

https://cloud.google.com/compute/docs/instance-groups/updating-migs#selective_updates no more explications neded

upvoted 1 times

jawulya 2 years, 11 months ago

The key here is "you don't want to update any running instances". Only opportunistic support this.

upvoted 1 times

heretolearnazure 2 years, 11 months ago

D is correct

upvoted 1 times

abhi52 3 years ago

Are the people creating these tests retarded? C is correct? how? Correct answer is D

upvoted 1 times

bhinar 3 years ago

If you want to apply updates automatically, set the type to proactive.

upvoted 1 times

BigfootPanda 3 years, 1 month ago

As we want to replace new instance, the update mode should be opportunistic. The default proactive mode would replace instance

upvoted 1 times

Question #117 *Topic 1*

Your company is designing its application landscape on Compute Engine. Whenever a zonal outage occurs, the application should be restored in another zone as quickly as possible with the latest application data. You need to design the solution to meet this requirement. What should you do?

A. Create a snapshot schedule for the disk containing the application data. Whenever a zonal outage occurs, use the latest snapshot to restore the disk in the same zone.

B. Configure the Compute Engine instances with an instance template for the application, and use a regional persistent disk for the application data. Whenever a zonal outage occurs, use the instance template to spin up the application in another zone in the same region. Use the regional persistent disk for the application data.

C. Create a snapshot schedule for the disk containing the application data. Whenever a zonal outage occurs, use the latest snapshot to restore the disk in another zone within the same region.

D. Configure the Compute Engine instances with an instance template for the application, and use a regional persistent disk for the application data. Whenever a zonal outage occurs, use the instance template to spin up the application in another region. Use the regional persistent disk for the application data.

Suggested Answer: *D*
Official Answer: D

Community vote distribution
B (97%) (3

TotoroChina Highly Voted 4 years, 7 months ago
Answer is B, it only request zonal resiliency. Regional persistent disk is a storage option that provides synchronous replication of data between two zones in a region. Regional persistent disks can be a good building block to use when you implement HA services in Compute Engine. https://cloud.google.com/compute/docs/disks/high-availability-regional-persistent-diskB. Agree, clearly it's B. Focus on keyword "zone"
upvoted 53 times

TotoroChina Highly Voted 4 years, 7 months ago
Answer is B, it only request zonal resiliency. Regional persistent disk is a storage option that provides synchronous replication of data between two zones in a region. Regional persistent disks can be a good building block to use when you implement HA services in Compute Engine.

https://cloud.google.com/compute/docs/disks/high-availability-regional-persistent-disk
upvoted 53 times

AmitRBS 3 years, 8 months ago
B. Agree, clearly it's B. Focus on keyword "zone"
upvoted 4 times

Ssoumya 4 years, 7 months ago
Answer is B
upvoted 14 times

Pime13 Most Recent 6 months ago
Regional Persistent Disks (Zero RPO): These disks provide synchronous replication of data between two zones within the same region. If one zone fails, the disk remains available in the secondary zone with all data intact up to the last write, ensuring you always have the latest application data. Instance Templates (Low RTO): An instance template stores the configuration of your VM (machine type, image, network settings). In a disaster scenario, you can use this template to immediately "spin up" a new instance in a healthy zone within seconds.
upvoted 1 times

plumbig11 1 year, 7 months ago
to be quickly as possible just change the zone in the same region configuring a; Compute Engine instance;
upvoted 1 times

oscarcampos 1 year, 9 months ago
why D ?
upvoted 1 times

mesodan 1 year, 11 months ago

It is B. As for D: Spinning up the application in another region might be too geographically distant, leading to higher latency and potential issues.
upvoted 1 times

hzaoui 2 years ago
A regional persistent disk is designed to provide synchronous replication of data between two zones in the same region, ensuring that data remains available even if one zone is affected by an outage. By using an instance template along with a regional disk, you can quickly create new instances in an available zone during a zonal outage and attach the regional persistent disk to continue operations with the latest application data.
upvoted 1 times

SSS987 2 years, 1 month ago
Can someone explain why not C - snapshot?
upvoted 1 times

SSS987 2 years, 1 month ago
Can someone explain why not C - snapshot?B - uses instance template which is ready for deploy. Option C requires manual configuration and this may take more time . But we need as quickly as possible .
upvoted 1 times

xaqanik 2 years ago
B - uses instance template which is ready for deploy. Option C requires manual configuration and this may take more time . But we need as quickly as possible .
upvoted 2 times

thewalker 2 years, 3 months ago
Selected Answer: D
Option D suggests using the same approach as Option B but restoring the application in another region instead of the same region. This approach provides high availability and disaster recovery across regions, making it suitable for applications that require high RTOs and minimal data loss.
upvoted 1 times

thewalker 2 years, 3 months ago
Option D suggests using the same approach as Option B but restoring the application in another region instead of the same region. This approach provides high availability and disaster recovery across regions, making it suitable for applications that require high RTOs and minimal data loss.can you use the regional persistent disk in a different region?apparently, nope.
upvoted 1 times

convers39 2 years, 1 month ago
can you use the regional persistent disk in a different region?
upvoted 3 times

convers39 2 years, 1 month ago
can you use the regional persistent disk in a different region?apparently, nope.
upvoted 3 times

don_v 2 years, 1 month ago
apparently, nope.
upvoted 1 times

JPA210 2 years, 4 months ago
Of course it is answer B. I would like to understand who chooses the right answers in examtopics! Choosing here option D is completely wrong. This takes people less instructed to be mistaken.
upvoted 2 times

AdityaGupta 2 years, 4 months ago
B. Configure the Compute Engine instances with an instance template for the application, and use a regional persistent disk for the application data. Whenever a zonal outage occurs, use the instance template to spin up the application in another zone in the same region. Use the regional persistent disk for the application data. Most Voted Why to change the region as mentioned in Option D, when the ask is different zone only.
upvoted 1 times

duzapo 2 years, 5 months ago
Answer is B, it only request zonal resiliency
upvoted 1 times

heretolearnazure 2 years, 5 months ago
It says restore in a zonal. Hence answer is B.
upvoted 1 times

mifrah 2 years, 10 months ago
B. In my opinion, I cannot use a regional disk in a different region!!! So, it can only be another zone in the same region. Therefore D must be wrong!
upvoted 2 times

SambuSoni 3 years ago

Answer B is Correct - since it talks about spin up application in different zone but same region. Whereas,D is incoorect , since its talking about spin up application in different region which is not our requirement.

upvoted 1 times

CosminCiuc 3 years ago

If it is a regional persistent disk created in region A for example. If I start the compute engine instance in the region B, how am I going to use a regional disk from region A (another region)? I do not think it is possible. So answer D should be excluded. I do believe that the correct answer is B.

upvoted 2 times

windsor_43 3 years, 1 month ago

The Answer is B Just had my exam today with a pass, this question was in the exam. Dated 31/12/22 Thanks to this site it was by far my most valuable

upvoted 5 times

Question #118 *Topic 1*

Your company has just acquired another company, and you have been asked to integrate their existing Google Cloud environment into your company's data center. Upon investigation, you discover that some of the RFC 1918 IP ranges being used in the new company's Virtual Private Cloud (VPC) overlap with your data center IP space. What should you do to enable connectivity and make sure that there are no routing conflicts when connectivity is established?

A. Create a Cloud VPN connection from the new VPC to the data center, create a Cloud Router, and apply new IP addresses so there is no overlapping IP space.

B. Create a Cloud VPN connection from the new VPC to the data center, and create a Cloud NAT instance to perform NAT on the overlapping IP space.

C. Create a Cloud VPN connection from the new VPC to the data center, create a Cloud Router, and apply a custom route advertisement to block the overlapping IP space.

D. Create a Cloud VPN connection from the new VPC to the data center, and apply a firewall rule that blocks the overlapping IP space.

Suggested Answer: *A*
Official Answer: A

Community vote distribution
A (50%) B (38%) C (12%)

VishalB Highly Voted 5 years ago
Correct Answer: A - IP Should not overlap so applying new IP address is the solutionA is not correct. "What should you do to enable connectivity and make sure that there are no routing conflicts when connectivity is established?" if you apply VPN con BGP, the actual IP address will be propagated to on prem environment with overlapping RFC1918 as result. B is correct with custom route
upvoted 44 times

VishalB Highly Voted 5 years ago
Correct Answer: A - IP Should not overlap so applying new IP address is the solution
upvoted 44 times

zanfo 4 years, 5 months ago
A is not correct. "What should you do to enable connectivity and make sure that there are no routing conflicts when connectivity is established?" if you apply VPN con BGP, the actual IP address will be propagated to on prem environment with overlapping RFC1918 as result. B is correct with custom route
upvoted 9 times

TotoroChina 5 years, 1 month ago
Answer is C. https://cloud.google.com/network-connectivity/docs/router/how-to/advertising-custom-ip
upvoted 37 times

TotoroChina 5 years, 1 month ago
Answer is C. https://cloud.google.com/network-connectivity/docs/router/how-to/advertising-custom-ipThe Q states to establish connectivity. This would merely prevent that. Ans is AI would also agree with C. Still, this part is confusing: "C. Create a Cloud VPN connection from the new VPC to the data center, create a Cloud Router, and apply a custom route advertisement to *block* the overlapping IP space." To *block*? Not to block. just to alias with advertised IP addresses.ANS is B https://cloud.google.com/architecture/best-practices-vpc-designYou can't use Cloud NAT according to this documentation: https://cloud.google.com/nat/docs/troubleshooting#overlapping-ip-addresses "Can I use Cloud NAT to connect a VPC network to another network to work around overlapping IP addresses? No, Cloud NAT cannot apply to any custom route whose next hop is not the default internet gateway. For example, Cloud NAT cannot apply to traffic sent to a next hop Cloud VPN tunnel, even if the destination is a publicly routable IP address."It will be a NAT Router instance, which will route the traffic. I have practically applied the configuration.B is NOT correct. Cloud NAT is specifically used for translating the IP address of the outbound packets destined to the Internet. But this question is about using VPN communication between two private IP address spaces (RFC1918). Cloud NAT cannot achieve the purpose here, you can't use Cloud NAT to translate from one private IP to another private ip. I would vote for C.Thanks for the clarification, just one question, without a solution like NAT or reip, the service on the devices with overlapping IP subnet will be unavailable for on-premise devices, not sure if the question also about this*you can't use Cloud NAT to translate from one private IP to another private ip to avoid overlapping ip range issue.Be careful because CloudNAT for private IPs will be GA soon
upvoted 37 times

meh009 4 years, 10 months ago

The Q states to establish connectivity. This would merely prevent that. Ans is A

upvoted 5 times

don_v 2 years, 7 months ago

I would also agree with C. Still, this part is confusing: "C. Create a Cloud VPN connection from the new VPC to the data center, create a Cloud Router, and apply a custom route advertisement to *block* the overlapping IP space." To *block*? Not to block. just to alias with advertised IP addresses.

upvoted 2 times

RKS_2021 5 years, 1 month ago

ANS is B https://cloud.google.com/architecture/best-practices-vpc-design

upvoted 10 times

RKS_2021 5 years, 1 month ago

ANS is B https://cloud.google.com/architecture/best-practices-vpc-designYou can't use Cloud NAT according to this documentation: https://cloud.google.com/nat/docs/troubleshooting#overlapping-ip-addresses "Can I use Cloud NAT to connect a VPC network to another network to work around overlapping IP addresses? No, Cloud NAT cannot apply to any custom route whose next hop is not the default internet gateway. For example, Cloud NAT cannot apply to traffic sent to a next hop Cloud VPN tunnel, even if the destination is a publicly routable IP address."It will be a NAT Router instance, which will route the traffic. I have practically applied the configuration.B is NOT correct. Cloud NAT is specifically used for translating the IP address of the outbound packets destined to the Internet. But this question is about using VPN communication between two private IP address spaces (RFC1918). Cloud NAT cannot achieve the purpose here, you can't use Cloud NAT to translate from one private IP to another private ip. I would vote for C.Thanks for the clarification, just one question, without a solution like NAT or reip, the service on the devices with overlapping IP subnet will be unavailable for on-premise devices, not sure if the question also about this*you can't use Cloud NAT to translate from one private IP to another private ip to avoid overlapping ip range issue.Be careful because CloudNAT for private IPs will be GA soon

upvoted 10 times

elenamatay 4 years, 7 months ago

You can't use Cloud NAT according to this documentation: https://cloud.google.com/nat/docs/troubleshooting#overlapping-ip-addresses "Can I use Cloud NAT to connect a VPC network to another network to work around overlapping IP addresses? No, Cloud NAT cannot apply to any custom route whose next hop is not the default internet gateway. For example, Cloud NAT cannot apply to traffic sent to a next hop Cloud VPN tunnel, even if the destination is a publicly routable IP address."

upvoted 17 times

RKS_2021 2 years, 10 months ago

It will be a NAT Router instance, which will route the traffic. I have practically applied the configuration.

upvoted 4 times

imgcp 5 years ago

B is NOT correct. Cloud NAT is specifically used for translating the IP address of the outbound packets destined to the Internet. But this question is about using VPN communication between two private IP address spaces (RFC1918). Cloud NAT cannot achieve the purpose here, you can't use Cloud NAT to translate from one private IP to another private ip. I would vote for C.Thanks for the clarification, just one question, without a solution like NAT or reip, the service on the devices with overlapping IP subnet will be unavailable for on-premise devices, not sure if the question also about this*you can't use Cloud NAT to translate from one private IP to another private ip to avoid overlapping ip range issue.Be careful because CloudNAT for private IPs will be GA soon

upvoted 13 times

imgcp 5 years ago

B is NOT correct. Cloud NAT is specifically used for translating the IP address of the outbound packets destined to the Internet. But this question is about using VPN communication between two private IP address spaces (RFC1918). Cloud NAT cannot achieve the purpose here, you can't use Cloud NAT to translate from one private IP to another private ip. I would vote for C.

upvoted 13 times

Bill831231 4 years, 10 months ago

Thanks for the clarification, just one question, without a solution like NAT or reip, the service on the devices with overlapping IP subnet will be unavailable for on-premise devices, not sure if the question also about this

upvoted 1 times

imgcp 5 years ago

*you can't use Cloud NAT to translate from one private IP to another private ip to avoid overlapping ip range issue.Be careful because CloudNAT for private IPs will be GA soon

upvoted 4 times

imgcp 5 years ago

*you can't use Cloud NAT to translate from one private IP to another private ip to avoid overlapping ip range issue.

upvoted 4 times

JohnnyBG 4 years, 1 month ago

Be careful because CloudNAT for private IPs will be GA soon

upvoted 1 times

0d8df8b Most Recent 2 months, 2 weeks ago

If RFC 1918 ranges overlap, routing becomes ambiguous: The network cannot reliably know whether traffic should stay local or go through the VPN. So the correct fix is to reassign IP ranges so the VPC and data center use non-overlapping CIDR ranges.

upvoted 1 times

Pime13 6 months ago

For successful routing across a VPN tunnel, IP ranges must be unique. If ranges overlap, routers cannot distinguish between local and remote traffic, leading to dropped packets or routing loops. Re-addressing one side of the connection ensures stable, long-term connectivity without the technical debt of complex translation layers.Option B (Cloud NAT): Traditionally, Cloud NAT (Public NAT) was only for outbound internet traffic and could not resolve overlapping private IP spaces across a VPN. While Google recently introduced Private NAT for this purpose, it adds significant management complexity compared to simple re-addressing and is often overkill for a standard merger integration. Option C (Custom Route Advertisements): Simply blocking the overlapping space via Cloud Router advertisements might stop routing conflicts, but it would also prevent communication with the resources in those specific ranges. It provides no actual connectivity for the overlapping segments.

upvoted 1 times

Pime13 6 months ago

Selected Answer: A

For successful routing across a VPN tunnel, IP ranges must be unique. If ranges overlap, routers cannot distinguish between local and remote traffic, leading to dropped packets or routing loops. Re-addressing one side of the connection ensures stable, long-term connectivity without the technical debt of complex translation layers.

upvoted 1 times

Pime13 6 months ago

Option B (Cloud NAT): Traditionally, Cloud NAT (Public NAT) was only for outbound internet traffic and could not resolve overlapping private IP spaces across a VPN. While Google recently introduced Private NAT for this purpose, it adds significant management complexity compared to simple re-addressing and is often overkill for a standard merger integration. Option C (Custom Route Advertisements): Simply blocking the overlapping space via Cloud Router advertisements might stop routing conflicts, but it would also prevent communication with the resources in those specific ranges. It provides no actual connectivity for the overlapping segments.

upvoted 1 times

05fe736 10 months, 2 weeks ago

The best approach is A. Create a Cloud VPN connection from the new VPC to the data center, create a Cloud Router, and apply new IP addresses so there is no overlapping IP space.

upvoted 2 times

CtrlX 11 months, 2 weeks ago

answer B. Was in my exam

upvoted 1 times

DM___1111 1 year, 1 month ago

This approach addresses the core problem of overlapping IP ranges: re-IP one of the networks (either the new company's VPC or your data center) so that there are no overlapping CIDR ranges. If IP ranges overlap, routers won't know which path to send traffic, leading to routing conflicts and connectivity failures. While Cloud NAT performs Network Address Translation, it's primarily designed for private instances to access the internet outbound. Using it to resolve overlapping internal IP ranges across a VPN tunnel adds significant complexity, management overhead, and can break certain types of connections. The best practice is to resolve the underlying IP overlap directly.

upvoted 2 times

Ankit267 1 year, 1 month ago

A is the best available option from the list, but in practice, you'd replace "apply new IP addresses" with implement a NAT gateway or proxy on one side to map the overlapping IP space to a non-overlapping range — preserving connectivity without re-IPing

upvoted 2 times

Zek 1 year, 5 months ago

https://cloud.google.com/nat/docs/overview#private-nat Assume that your Google Cloud resources in a VPC network need to communicate with destinations in a VPC, on-premises, or other cloud provider network that is owned by a different business unit. However, the destination network contains subnets whose IP addresses overlap with the IP addresses of your VPC network. In this scenario, you create a Cloud NAT gateway for Private NAT that translates traffic between the subnets in your VPC network to the non-overlapping subnets of the other network.

upvoted 3 times

Peto12 1 year, 7 months ago

With A you need to apply new IP addresses, with B you can use private NAT.

upvoted 1 times

andreacola 1 year, 9 months ago

Assume that the resources in your VPC network need to communicate with the resources in a VPC network or an on-premises or other cloud provider network that is owned by a different business unit. However, that network contains subnets whose IP addresses overlap with the IP addresses of your VPC network. In this scenario, you create a Private NAT gateway that translates traffic between the subnets in your VPC network to the non-overlapping subnets of the other network.

upvoted 5 times

Abhinavchawlac2d 1 year, 10 months ago

Correct Option: B. Create a Cloud VPN connection from the new VPC to the data center, and create a Cloud NAT instance to perform NAT on the overlapping IP space. This option effectively allows you to connect the two environments while addressing the overlapping IP space issue through NAT, ensuring that the VMs can communicate without conflicts.

upvoted 2 times

3fd692e 1 year, 10 months ago

There is a Private NAT you can use and is specifically designed to resolve overlapping private IP issues: https://medium.com/niveus-solutions/private-cloud-nat-and-why-we-need-it-on-gcp-f6ad0c96facb#:~:text=Private%20Cloud%20NAT%20with%20NCC,helps%20connect%20onprem%20to%20gcp.

upvoted 5 times

lucaluca1982 2 years ago

Given that you are not going out to the internet and you need to use a Cloud Router for your VPC, you need to ensure that there is no overlap in the IP ranges between your data center and the newly acquired company's VPC. The best approach to manage this without renumbering the entire network is to use Network Address Translation (NAT) to handle the overlapping IP addresses.

upvoted 1 times

nhatne 2 years, 1 month ago

would go for B

upvoted 2 times

Sephethus 2 years, 2 months ago

The answer is B. Cloud VPN and Cloud NAT help you get around this problem easily without all the work of creating a new subnet and reassigning IPs to everything. Cloud NAT: Network Address Translation (NAT) allows you to translate IP addresses in your VPC to a different IP range, avoiding conflicts with overlapping IP ranges in your data center. This ensures that traffic can flow between the environments without routing conflicts. Cloud VPN: Establishing a Cloud VPN connection provides secure connectivity between the new VPC and your data center. By combining this with Cloud NAT, you can effectively manage and resolve the IP address overlap.Cloud NAT does not directly resolve IP address conflicts due to overlapping ranges. Cloud NAT is typically used for instances without external IP addresses to access the internet while preserving their internal IPs for internal communications.

upvoted 1 times

Sephethus 2 years, 2 months ago

The answer is B. Cloud VPN and Cloud NAT help you get around this problem easily without all the work of creating a new subnet and reassigning IPs to everything.

Cloud NAT: Network Address Translation (NAT) allows you to translate IP addresses in your VPC to a different IP range, avoiding conflicts with overlapping IP ranges in your data center. This ensures that traffic can flow between the environments without routing conflicts. Cloud VPN: Establishing a Cloud VPN connection provides secure connectivity between the new VPC and your data center. By combining this with Cloud NAT, you can effectively manage and resolve the IP address overlap.

upvoted 1 times

Sephethus 2 years, 2 months ago

Cloud NAT does not directly resolve IP address conflicts due to overlapping ranges. Cloud NAT is typically used for instances without external IP addresses to access the internet while preserving their internal IPs for internal communications.

upvoted 1 times

eff12c1 2 years, 2 months ago

Selected Answer: B

Using Cloud NAT to translate overlapping IP addresses is the most effective solution to ensure seamless connectivity between the new company's VPC and your company's data center without routing conflicts. This approach avoids the complexity of reconfiguring IP addresses and ensures that both networks can communicate effectively. https://cloud.google.com/nat/docs/overview#private-nat

upvoted 4 times

eff12c1 2 years, 2 months ago

Using Cloud NAT to translate overlapping IP addresses is the most effective solution to ensure seamless connectivity between the new company's VPC and your company's data center without routing conflicts. This approach avoids the complexity of reconfiguring IP addresses and ensures that

both networks can communicate effectively. https://cloud.google.com/nat/docs/overview#private-natIt is not NAT, we are not going out to internet. We need cloud routerNAT doesn't necessarily need to be used only for internet communication

upvoted 4 times

ccpmad 2 years, 2 months ago

It is not NAT, we are not going out to internet. We need cloud router

upvoted 1 times

ccpmad 2 years, 2 months ago

It is not NAT, we are not going out to internet. We need cloud routerNAT doesn't necessarily need to be used only for internet communication

upvoted 1 times

VegasDegenerate 1 year, 7 months ago

NAT doesn't necessarily need to be used only for internet communication

upvoted 1 times

Question #119 *Topic 1*

You need to migrate Hadoop jobs for your company's Data Science team without modifying the underlying infrastructure. You want to minimize costs and infrastructure management effort. What should you do?

A. Create a Dataproc cluster using standard worker instances.

B. Create a Dataproc cluster using preemptible worker instances.

C. Manually deploy a Hadoop cluster on Compute Engine using standard instances.

D. Manually deploy a Hadoop cluster on Compute Engine using preemptible instances.

Suggested Answer: *A*

Reference: https://cloud.google.com/architecture/hadoop/hadoop-gcp-migration-jobs

Community vote distribution

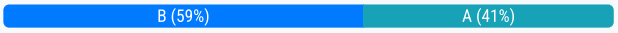

TotoroChina Highly Voted 5 years, 1 month ago

Should be B, you want to minimize costs. https://cloud.google.com/dataproc/docs/concepts/compute/secondary-vms#preemptible_and_non-preemptible_secondary_workersHi TotoroChina, I had the same thought when I first read the question - the problem I see is, in real business I think you would try to mix preemtible instances and on-demand instances... Here you have to choose between only preemtible instances and on-demand instances... Preemptible instances have some downsides - so we would need more details and ideally a mixed approach. That's why both answers might be correcy, a) and b)... Do you see that different? Thanks! Cheers, D.but you need to reduce management overhead so B if you create a cluster manually and create and maintain GCE is not the way to goB requires to create new instances at least every 24h."without modifying the underlying infrastructure" is the watch word. Most likely did not utilize preemptible on-premisesA cost-savings consideration: Using preemptible VMs does not always save costs since preemptions can cause longer job execution with resulting higher job costs. This is mentioned in above link So I think Ans should be AI think it's A. The question does not mention anything about minimize the costs, all the questions in GCP exams that require minimize the costs as requirement literally mention that in the question. Also in order to minimize the costs you need to build jobs that are fault tolerant, as workers instances are preemptible. This also requires some kind of Dev investment of work. So if not mentioned in the question fault tolerant and minimize costs then is not required/needed. Doc states below: Only use preemptible nodes for jobs that are fault-tolerant or that are low enough priority that occasional job failure won't disrupt your business.The question is very short and literally says "You want to minimize costs" lollearn to read before apply to the examIt's dissapointing that you waste your time writting on this topic instead of payin

upvoted 69 times

TotoroChina Highly Voted 5 years, 1 month ago

Should be B, you want to minimize costs. https://cloud.google.com/dataproc/docs/concepts/compute/secondary-vms#preemptible_and_non-preemptible_secondary_workers

upvoted 69 times

XDevX 5 years, 1 month ago

Hi TotoroChina, I had the same thought when I first read the question - the problem I see is, in real business I think you would try to mix preemtible instances and on-demand instances... Here you have to choose between only preemtible instances and on-demand instances... Preemptible instances have some downsides - so we would need more details and ideally a mixed approach. That's why both answers might be correcy, a) and b)... Do you see that different? Thanks! Cheers, D.

upvoted 5 times

XDevX 5 years, 1 month ago

Hi TotoroChina, I had the same thought when I first read the question - the problem I see is, in real business I think you would try to mix preemtible instances and on-demand instances... Here you have to choose between only preemtible instances and on-demand instances... Preemptible instances have some downsides - so we would need more details and ideally a mixed approach. That's why both answers might be correcy, a) and b)... Do you see that different? Thanks! Cheers, D.but you need to reduce management overhead so B if you create a cluster manually and create and maintain GCE is not the way to goB requires to create new instances at least every 24h.

upvoted 5 times

kopper2019 5 years, 1 month ago

but you need to reduce management overhead so B if you create a cluster manually and create and maintain GCE is not the way to go

upvoted 5 times

kopper2019 5 years, 1 month ago

but you need to reduce management overhead so B if you create a cluster manually and create and maintain GCE is not the way to goB requires to create new instances at least every 24h.

upvoted 5 times

HenkH 3 years, 9 months ago

B requires to create new instances at least every 24h.

upvoted 2 times

Sukon_Desknot 3 years, 9 months ago

"without modifying the underlying infrastructure" is the watch word. Most likely did not utilize preemptible on-premises

upvoted 8 times

Yogi42 3 years, 7 months ago

A cost-savings consideration: Using preemptible VMs does not always save costs since preemptions can cause longer job execution with resulting higher job costs. This is mentioned in above link So I think Ans should be A

upvoted 3 times

ale_brd_111 3 years, 8 months ago

I think it's A. The question does not mention anything about minimize the costs, all the questions in GCP exams that require minimize the costs as requirement literally mention that in the question. Also in order to minimize the costs you need to build jobs that are fault tolerant, as workers instances are preemptible. This also requires some kind of Dev investment of work. So if not mentioned in the question fault tolerant and minimize costs then is not required/needed. Doc states below: Only use preemptible nodes for jobs that are fault-tolerant or that are low enough priority that occasional job failure won't disrupt your business.The question is very short and literally says "You want to minimize costs" lollearn to read before apply to the examIt's dissapointing that you waste your time writting on this topic instead of paying attention at the questions.

upvoted 8 times

ale_brd_111 3 years, 8 months ago

I think it's A. The question does not mention anything about minimize the costs, all the questions in GCP exams that require minimize the costs as requirement literally mention that in the question. Also in order to minimize the costs you need to build jobs that are fault tolerant, as workers instances are preemptible. This also requires some kind of Dev investment of work. So if not mentioned in the question fault tolerant and minimize costs then is not required/needed.

Doc states below: Only use preemptible nodes for jobs that are fault-tolerant or that are low enough priority that occasional job failure won't disrupt your business.

upvoted 8 times

jasenmornin 3 years, 8 months ago

The question is very short and literally says "You want to minimize costs" lol

upvoted 14 times

xprtz1 3 years, 8 months ago

learn to read before apply to the exam

upvoted 8 times

zetalexg 3 years, 8 months ago

It's dissapointing that you waste your time writting on this topic instead of paying attention at the questions.

upvoted 30 times

firecloud 5 years ago

It's A, the primary workers can only be standard, where secondary workers can be preemtible.------In addition to using standard Compute Engine VMs as Dataproc workers (called "primary" workers), Dataproc clusters can use "secondary" workers. There are two types of secondary workers: preemptible and non-preemptible. All secondary workers in your cluster must be of the same type, either preemptible or non-preemptible. The default is preemptible.

upvoted 35 times

firecloud 5 years ago

It's A, the primary workers can only be standard, where secondary workers can be preemtible.------In addition to using standard Compute Engine VMs as Dataproc workers (called "primary" workers), Dataproc clusters can use "secondary" workers. There are two types of secondary workers: preemptible and non-preemptible. All secondary workers in your cluster must be of the same type, either preemptible or non-preemptible. The default is preemptible.agreed

upvoted 35 times

Manh 4 years, 11 months ago

agreed

upvoted 2 times

Pime13 Most Recent 6 months ago

minimize costs

upvoted 1 times

CtrlX 11 months, 2 weeks ago

answer B. Was in my exam

upvoted 1 times

david_tay 1 year, 5 months ago

it mentioned cost saving, hence B is fine as Hadoop is usually used for batch processing so a bit of downtime is tolerable.

upvoted 2 times

3a7557a 1 year, 6 months ago

A. Create a Dataproc cluster using standard worker instances. you can use premtible instnace however should be less then 50% totaak wroker instances... could impact staiblirty... cons transient task failures. follow best practice: General Recommendations: Start Conservatively: Begin with a smaller percentage of preemptible workers (e.g., less than 30% of the total worker nodes). Monitor your jobs closely to see how they perform and adjust as needed. Don't Exceed 50%: Google generally recommends keeping the number of preemptible workers below 50% of the total worker nodes in your cluster. 1 This helps maintain stability and reduces the risk of significant job disruptions.

upvoted 1 times

plumbig11 1 year, 7 months ago

Dataproc, as they want to save costs, preemptible is the best option, for sure.

upvoted 2 times

Lestrang 1 year, 10 months ago

Probably A - Primary workers must be standard - Preemptible doesn't always save cost - Infrastructure on prem doesn't have spot machines - You cannot choose spot/preemptible when creating the cluster, only when provisioning secondary nodes, which are actually preemptible by default. - They do not store data, only do data processing

upvoted 1 times

pcamaster 1 year, 10 months ago

B: being a multiple-choice question, we need to focus on explicit keywords here. Management effort => Managed service -> Dataproc. Cost-optimization => Preemptible. For ones who say "but that also requires fault toleration": well, there is no explicit keyword in the question says "we have critical jobs" or "out data scientists team has not takes into account toleration". So we must not assume that's needed.

upvoted 2 times

afsarkhan 2 years, 1 month ago

I will go with A , reason preemptible instances are unpredictable and there is no mention of work criticallity. So my answer is A against B

upvoted 1 times

Gino17m 2 years, 3 months ago

A - only secondary workers can be preemptible and "Using preemptible VMs does not always save costs since preemptions can cause longer job execution with resulting higher job costs" according to: https://cloud.google.com/dataproc/docs/concepts/compute/secondary-vms#preemptible_and_non-preemptible_secondary_workers

upvoted 1 times

dija123 2 years, 4 months ago

Agree with B

upvoted 2 times

Diwz 2 years, 4 months ago

Answer is B. The secondary worker type instance for default Dataproc cluster is preemptible VMs.
https://cloud.google.com/dataproc/docs/concepts/compute/secondary-vms

upvoted 1 times

shashii82 2 years, 5 months ago

Dataproc: Dataproc is a fully managed Apache Spark and Hadoop service on Google Cloud Platform. It allows you to run clusters without the need to manually deploy and manage Hadoop clusters on Compute Engine. Preemptible Worker Instances: Preemptible instances are short-lived, cost-effective virtual machine instances that are suitable for fault-tolerant and batch processing workloads. Since Hadoop jobs can often tolerate interruptions, using preemptible instances can significantly reduce costs. Option B leverages the benefits of Dataproc for managing Hadoop clusters without the need for manual deployment and takes advantage of preemptible instances to minimize costs. This aligns well with the goal of minimizing both costs and infrastructure management efforts.

upvoted 1 times

VidhyaBupesh 2 years, 5 months ago

Using preemptible VMs does not always save costs since preemptions can cause longer job execution with resulting higher job costs

upvoted 1 times

Amrita2012 2 years, 6 months ago

Using standard Compute Engine VMs as Dataproc workers (called "primary" workers), Preemptible can be only used for secondary workers hence A is valid answer

upvoted 1 times

d0094d6 2 years, 6 months ago

You want to minimize costs and infrastructure management effort > B

upvoted 3 times

Question #120 *Topic 1*

Your company has a project in Google Cloud with three Virtual Private Clouds (VPCs). There is a Compute Engine instance on each VPC. Network subnets do not overlap and must remain separated. The network configuration is shown below.
Instance #1 is an exception and must communicate directly with both Instance #2 and Instance #3 via internal IPs. How should you accomplish this?

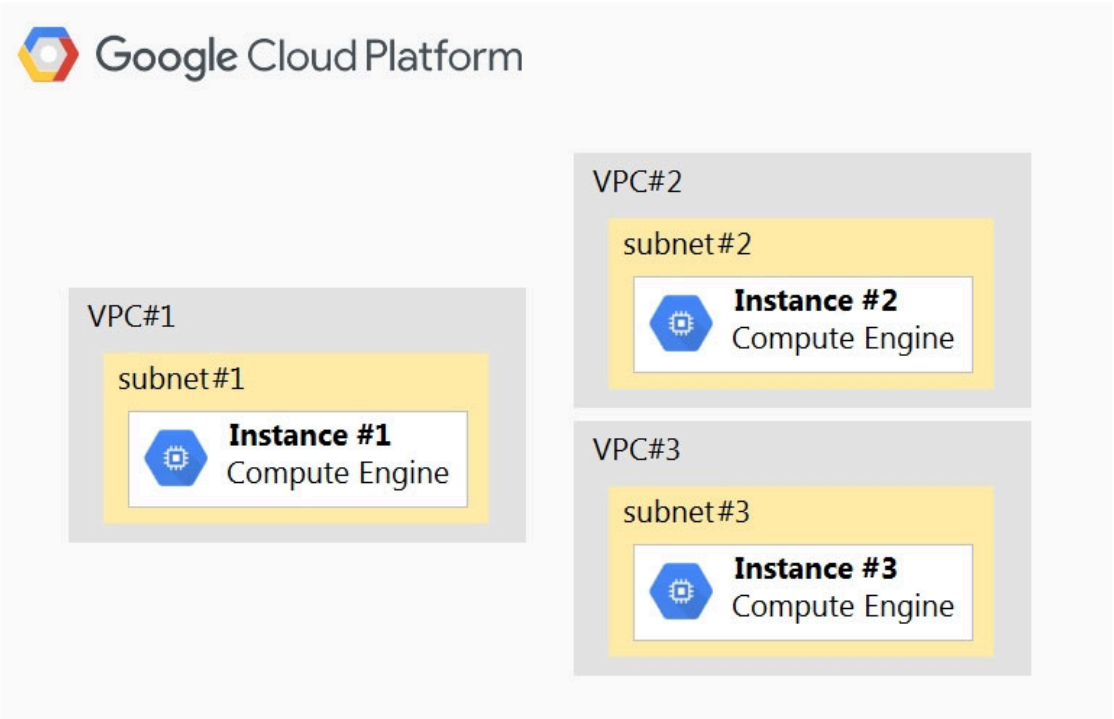

A. Create a cloud router to advertise subnet #2 and subnet #3 to subnet #1.

B. Add two additional NICs to Instance #1 with the following configuration: ג€¢ NIC1 ג–‹ VPC: VPC #2 ג–‹ SUBNETWORK: subnet #2 ג€¢ NIC2 ג–‹ VPC: VPC #3 ג–‹ SUBNETWORK: subnet #3 Update firewall rules to enable traffic between instances.

C. Create two VPN tunnels via CloudVPN: 1 ¢€ג between VPC #1 and VPC #2. 1 ¢€ג between VPC #2 and VPC #3. Update firewall rules to enable traffic between the instances.

D. Peer all three VPCs: ג€¢ Peer VPC #1 with VPC #2. ג€¢ Peer VPC #2 with VPC #3. Update firewall rules to enable traffic between the instances.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (74%) D (13%) C (13%)

XDevX Highly Voted 5 years, 1 month ago
According to my understanding the requirement is that only VM1 shall be able to communicate with VM2 and VM3, but not VM2 with VM3. We can exclude d) as d) would enable VM2 to communicate with VM3 as well - my assumption is, that if the quizzer wanted that d) is the correct answer, he would make just 2 peerings - 1x between VM1 and VM2 and 1x between VM1 and VM3 repectively the VPCs. We can exclude c) as well - there is no connection between VPC1 and VPC3. IMHO a) will not work. So the only correct answer seems to be b) - what I don't understand is why we have to update the firewall rules as IMHO the default firewall rules enable such communication (maybe some restrictive rules are implemented - not enough details in the question to clarify that part). Please correct me if I am wrong.Correct, maybe fw on the VMThe answer is "B". The following link has this - "Use multiple network interfaces when an individual instance needs access to more than one VPC network, but you don't want to connect both networks directly." https://cloud.google.com/vpc/docs/multiple-interfaces-conceptsB will not work. VM instances within a VPC network can communicate among themselves using internal IP addresses as long as firewall rules permit. However, no internal IP address communication is allowed between networks, unless you set up mechanisms such as VPC Network Peering or Cloud VPN.this link says VM can have multiple NICs and attached to different VPCs. https://cloud.google.com/vpc/docs/create-use-multiple-interfaces so B is the answerI think it is because the instances

are in separate VPCs. "Google Cloud Virtual Private Cloud (VPC) networks are by default isolated private networking domains. Networks have a global scope and contain regional subnets. VM instances within a VPC network can communicate among themselves using internal IP addresses as long as firewall rules permit. However, NO INTERNAL IP ADDRESS COMMUNICATION IS ALLOWED BETWEEN networks, unless you set up mechanisms such as

upvoted 25 times

XDevX Highly Voted 5 years, 1 month ago

According to my understanding the requirement is that only VM1 shall be able to communicate with VM2 and VM3, but not VM2 with VM3. We can exclude d) as d) would enable VM2 to communicate with VM3 as well - my assumption is, that if the quizzer wanted that d) is the correct answer, he would make just 2 peerings - 1x between VM1 and VM2 and 1x between VM1 and VM3 repectively the VPCs. We can exclude c) as well - there is no connection between VPC1 and VPC3. IMHO a) will not work. So the only correct answer seems to be b) - what I don't understand is why we have to update the firewall rules as IMHO the default firewall rules enable such communication (maybe some restrictive rules are implemented - not enough details in the question to clarify that part). Please correct me if I am wrong.

upvoted 25 times

JeffClarke111 5 years, 1 month ago

Correct, maybe fw on the VM

upvoted 2 times

Ishu_awsguy 3 years, 12 months ago

The answer is "B". The following link has this - "Use multiple network interfaces when an individual instance needs access to more than one VPC network, but you don't want to connect both networks directly." https://cloud.google.com/vpc/docs/multiple-interfaces-concepts

upvoted 13 times

Ishu_awsguy 3 years, 12 months ago

The answer is "B". The following link has this - "Use multiple network interfaces when an individual instance needs access to more than one VPC network, but you don't want to connect both networks directly." https://cloud.google.com/vpc/docs/multiple-interfaces-conceptsB will not work. VM instances within a VPC network can communicate among themselves using internal IP addresses as long as firewall rules permit. However, no internal IP address communication is allowed between networks, unless you set up mechanisms such as VPC Network Peering or Cloud VPN.this link says VM can have multiple NICs and attached to different VPCs. https://cloud.google.com/vpc/docs/create-use-multiple-interfaces so B is the answer

upvoted 13 times

Pankaj_007 3 years, 9 months ago

B will not work. VM instances within a VPC network can communicate among themselves using internal IP addresses as long as firewall rules permit. However, no internal IP address communication is allowed between networks, unless you set up mechanisms such as VPC Network Peering or Cloud VPN.this link says VM can have multiple NICs and attached to different VPCs. https://cloud.google.com/vpc/docs/create-use-multiple-interfaces so B is the answer

upvoted 1 times

Pankaj_007 3 years, 9 months ago

B will not work. VM instances within a VPC network can communicate among themselves using internal IP addresses as long as firewall rules permit. However, no internal IP address communication is allowed between networks, unless you set up mechanisms such as VPC Network Peering or Cloud VPN.

upvoted 1 times

sameer2803 3 years, 7 months ago

this link says VM can have multiple NICs and attached to different VPCs. https://cloud.google.com/vpc/docs/create-use-multiple-interfaces so B is the answer

upvoted 6 times

lazybeanbag 5 years ago

I think it is because the instances are in separate VPCs.

"Google Cloud Virtual Private Cloud (VPC) networks are by default isolated private networking domains. Networks have a global scope and contain regional subnets. VM instances within a VPC network can communicate among themselves using internal IP addresses as long as firewall rules permit. However, NO INTERNAL IP ADDRESS COMMUNICATION IS ALLOWED BETWEEN networks, unless you set up mechanisms such as VPC Network Peering or Cloud VPN."

The instructions for setting up multiple interfaces tells you to check your firewall rules as as the firewall rules of the VPC apply to the network interface that it is attached to.

https://cloud.google.com/vpc/docs/multiple-interfaces-concepts#firewall_rules_and_multiple_network_interfaces

upvoted 8 times

lazybeanbag 5 years ago

I think it is because the instances are in separate VPCs. "Google Cloud Virtual Private Cloud (VPC) networks are by default isolated private networking domains. Networks have a global scope and contain regional subnets. VM instances within a VPC network can communicate among themselves using internal IP addresses as long as firewall rules permit. However, NO INTERNAL IP ADDRESS COMMUNICATION IS ALLOWED BETWEEN networks, unless you set up mechanisms such as VPC Network Peering or Cloud VPN." The instructions for setting up multiple interfaces tells you to check your firewall rules as as the firewall rules of the VPC apply to the network interface that it is attached to. https://cloud.google.com/vpc/docs/multiple-interfaces-concepts#firewall_rules_and_multiple_network_interfacesThe answer is "B". The following link has this - "Use multiple network interfaces when an individual instance needs access to more than one VPC network, but you don't want to connect both networks directly." https://cloud.google.com/vpc/docs/multiple-interfaces-conceptsyou can not add additional network interface to existing VM's

upvoted 8 times

lshu_awsguy 3 years, 12 months ago

The answer is "B". The following link has this - "Use multiple network interfaces when an individual instance needs access to more than one VPC network, but you don't want to connect both networks directly." https://cloud.google.com/vpc/docs/multiple-interfaces-conceptsyou can not add additional network interface to existing VM's

upvoted 5 times

b6f53d8 2 years, 10 months ago

you can not add additional network interface to existing VM's

upvoted 1 times

MamthaSJ 5 years, 1 month ago

Answer is B

upvoted 11 times

MamthaSJ 5 years, 1 month ago

Answer is BInstances are exist. You can not add or remove additional NICs to a VM

upvoted 11 times

coutcin 4 years, 3 months ago

Instances are exist. You can not add or remove additional NICs to a VM

upvoted 7 times

plumbig11 Most Recent 1 year, 7 months ago

Add two additional NICs to Instance

upvoted 2 times

awsgcparch 2 years ago

Direct Connectivity: Adding multiple NICs to Instance #1 allows it to be part of multiple VPCs directly. This configuration enables direct communication with Instance #2 and Instance #3 via internal IPs without requiring additional routing configurations. Simplicity: This approach is straightforward and avoids the complexity of setting up VPC peering or VPN tunnels. It ensures that only Instance #1 has access to both VPC #2 and VPC #3, maintaining the separation of the other VPCs.

upvoted 2 times

afsarkhan 2 years, 1 month ago

VPC peering will allow access to instance 2 & 3 from 1 with internal IP, with necessary firewall rules added.

upvoted 1 times

dija123 2 years, 4 months ago

B for sure

upvoted 2 times

shashii82 2 years, 5 months ago

Option B allows you to add additional NICs to Instance #1, each connected to a different VPC, facilitating direct communication between Instance #1 and the other instances while maintaining separate subnets.

upvoted 1 times

kshlgpt 2 years, 7 months ago

B is wrong. NIC can only be configured while creating the instance. Here the instance is already created. C is correct answer. Refer limitation in this link: https://cloud.google.com/vpc/docs/create-use-multiple-interfaces

upvoted 2 times

AdityaGupta 2 years, 10 months ago

Router, VPN and VPC Peering for all 3 network is not required. Only option B solves the given scenario.

upvoted 2 times