GETCERTKEY

GETCERTKEY
100% guarantee you pass IT cert exam!

http://www.getcertkey.com

No help, Full refund!

Exam : **Associate-Cloud-Engineer**

Title : Google Associate Cloud Engineer Exam

Vendor : Google

Version : DEMO

NO.1 You have been asked to set up Object Lifecycle Management for objects stored in storage buckets. The objects are written once and accessed frequently for 30 days. After 30 days, the objects are not read again unless there is a special need. The object should be kept for three years, and you need to minimize cost.
What should you do?
A. Set up a policy that uses Nearline storage for 30 days and then moves to Archive storage for three years.
B. Set up a policy that uses Standard storage for 30 days and then moves to Archive storage for three years.
C. Set up a policy that uses Nearline storage for 30 days, then moves the Coldline for one year, and then moves to Archive storage for two years.
D. Set up a policy that uses Standard storage for 30 days, then moves to Coldline for one year, and then moves to Archive storage for two years.
Answer: B
Explanation:
The key to understand the requirement is: "The objects are written once and accessed frequently for 30 days" Standard Storage Standard Storage is best for data that is frequently accessed ("hot" data) and/or stored for only brief periods of time.
Archive Storage
Archive Storage is the lowest-cost, highly durable storage service for data archiving, online backup, and disaster recovery. Unlike the "coldest" storage services offered by other Cloud providers, your data is available within milliseconds, not hours or days. Archive Storage is the best choice for data that you plan to access less than once a year.
https://cloud.google.com/storage/docs/storage-classes#standard

NO.2 You are running multiple microservices in a Kubernetes Engine cluster. One microservice is rendering images. The microservice responsible for the image rendering requires a large amount of CPU time compared to the memory it requires. The other microservices are workloads that are optimized for n1-standard machine types. You need to optimize your cluster so that all workloads are using resources as efficiently as possible.
What should you do?
A. Assign the pods of the image rendering microservice a higher pod priority than the older microservices
B. Create a node pool with compute-optimized machine type nodes for the image rendering microservice.
Use the node pool with general-purpose machine type nodes for the other microservices
C. Use the node pool with general-purpose machine type nodes for lite mage rendering microservice .
Create a nodepool with compute-optimized machine type nodes for the other microservices
D. Configure the required amount of CPU and memory in the resource requests specification of the image rendering microservice deployment.
Keep the resource requests for the other microservices at the default
Answer: B

NO.3 You want to host your video encoding software on Compute Engine. Your user base is growing

rapidly, and users need to be able to encode their videos at any time without interruption or CPU limitations. You must ensure that your encoding solution is highly available, and you want to follow Google-recommended practices to automate operations. What should you do?

A. Deploy your solution on multiple standalone Compute Engine instances, and increase the number of existing instances when CPU utilization on Cloud Monitoring reaches a certain threshold.

B. Deploy your solution on multiple standalone Compute Engine instances, and replace existing instances with high-CPU instances when CPU utilization on Cloud Monitoring reaches a certain threshold.

C. Deploy your solution to an instance group, and increase the number of available instances whenever you see high CPU utilization in Cloud Monitoring.

D. Deploy your solution to an instance group, and set the autoscaling based on CPU utilization.

Answer: D

Explanation:

You can create a managed instance group with autoscaling enabled based on CPU utilization.
This way appropriate number of instances can be added or removed based on the CPU metrics.
https://cloud.google.com/compute/docs/instance-groups/create-mig-with-basic-autoscaling
https://cloud.google.com/sdk/gcloud/reference/compute/instance-groups/managed/set- autoscalin g

NO.4 You have a website hosted on App Engine standard environment. You want 1% of your users to see a new test version of the website. You want to minimize complexity. What should you do?

A. Deploy the new version in the same application and use the --migrate option.

B. Deploy the new version in the same application and use the --splits option to give a weight of 99 to the current version and a weight of 1 to the new version.

C. Create a new App Engine application in the same project. Deploy the new version in that application.
Use the App Engine library to proxy 1% of the requests to the new version.

D. Create a new App Engine application in the same project. Deploy the new version in that application.
Configure your network load balancer to send 1% of the traffic to that new application.

Answer: B

Explanation:

There is only one app engine can be created per project.
https://cloud.google.com/appengine/docs/standard/python/splitting-traffic#gcloud

NO.5 You want to configure 10 Compute Engine instances for availability when maintenance occurs. Your requirements state that these instances should attempt to automatically restart if they crash. Also, the instances should be highly available including during system maintenance. What should you do?

A. Create an instance template for the instances.
Set the `Automatic Restart' to on. Set the `On-host maintenance' to Migrate VM instance.
Add the instance template to an intsance group.

B. Create an instance template for the instances.
`Automatic Restart' to off. Set `On-host maintenance' to Terminate VM instances.
Add the instance template to an instance group.

C. Create an instance group for the instances.
Set the `Autohealing' health check to healthy (HTTP).
D. Create an instance group for the instance.
Verify that the `Advanced creation options' setting for `do not retry machine creation' is set to off.
Answer: A
Explanation:
onHostMaintenance: Determines the behavior when a maintenance event occurs that might cause your instance to reboot.
[Default] MIGRATE, which causes Compute Engine to live migrate an instance when there is a maintenance event.
TERMINATE, which stops an instance instead of migrating it.
automaticRestart: Determines the behavior when an instance crashes or is stopped by the system.
[Default] true, so Compute Engine restarts an instance if the instance crashes or is stopped.
false, so Compute Engine does not restart an instance if the instance crashes or is stopped.
https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options

NO.6 You have developed an application that consists of multiple microservices, with each microservice packaged in its own Docker container image. You want to deploy the entire application on Google Kubernetes Engine so that each microservice can be scaled individually.
What should you do?
A. Create and deploy a Custom Resource Definition per microservice.
B. Create and deploy a Docker Compose File.
C. Create and deploy a Job per microservice.
D. Create and deploy a Deployment per microservice.
Answer: D
Explanation:
Can deploy each service through
kubectl apply -f <deployment_config.yaml>

NO.7 The sales team has a project named Sales Data Digest that has the ID acme-data-digest You need to set up similar Google Cloud resources for the marketing team but their resources must be organized independently of the sales team. What should you do?
A. Grant the Project Editor role to the Marketing learn for acme data digest
B. Create a Project Lien on acme-data digest and then grant the Project Editor role to the Marketing team
C. Create another protect with the ID acme-marketing-data-digest for the Marketing team and deploy the resources there
D. Create a new protect named Meeting Data Digest and use the ID acme-data-digest Grant the Project Editor role to the Marketing team.
Answer: C
Explanation:
The resources for the marketing team should be independent from the Sales team. Resources are tied and separated by projects.

NO.8 You have one GCP account running in your default region and zone and another account

running in a non-default region and zone. You want to start a new Compute Engine instance in these two Google Cloud Platform accounts using the command line interface. What should you do?

A. Create two configurations using gcloud config configurations create [NAME].
Run gcloud config configurations activate [NAME]to switch between accounts when running the commands to start the Compute Engine instances.

B. Create two configurations using gcloud config configurations create [NAME].
Run gcloud configurations listto start the Compute Engine instances.

C. Activate two configurations using gcloud configurations activate [NAME].
Run gcloud config listto start the Compute Engine instances.

D. Activate two configurations using gcloud configurations activate [NAME].
Run gcloud configurations listto start the Compute Engine instances.

Answer: A
Explanation:
You can create different configurations for each account and create compute instances in each account by activating the respective account.

NO.9 You need to create a Compute Engine instance in a new project that doesn't exist yet. What should you do?

A. Using the Cloud SDK, create a new project, enable the Compute Engine API in that project, and then create the instance specifying your new project.

B. Enable the Compute Engine API in the Cloud Console, use the Cloud SDK to create the instance, and then use the -pproject flag to specify a new project.

C. Using the Cloud SDK, create the new instance, and use the -pproject flag to specify the new project.
Answer yes when prompted by Cloud SDK to enable the Compute Engine API.

D. Enable the Compute Engine API in the Cloud Console. Go to the Compute Engine section of the Console to create a new instance, and look for the Create In A New Project option in the creation form.

Answer: A
Explanation:
Before you begin:
1. In the Cloud Console, on the project selector page, select or create a Cloud project.
2. Make sure that billing is enabled for your Google Cloud project. Learn how to confirm billing is enabled for your project.
To use the gcloud command-line tool for this quickstart, you must first install and initialize the Cloud SDK:
1. Download and install the Cloud SDK using the instructions given on Installing Google Cloud SDK.
2. Initialize the SDK using the instructions given on Initializing Cloud SDK.
To use gcloud in Cloud Shell for this quickstart, first activate Cloud Shell using the instructions given on Starting Cloud Shell.
https://cloud.google.com/ai-platform/deep-learning-vm/docs/quickstart-cli#before-you-begin

NO.10 You have a single binary application that you want to run on Google Cloud Platform. You decided to automatically scale the application based on underlying infrastructure CPU usage. Your organizational policies require you to use virtual machines directly. You need to ensure that the

application scaling is operationally efficient and completed as quickly as possible. What should you do?

A. Create a Google Kubernetes Engine cluster, and use horizontal pod autoscaling to scale the application.

B. Create an instance template, and use the template in a managed instance group with autoscaling configured.

C. Create an instance template, and use the template in a managed instance group that scales up and down based on the time of day.

D. Use a set of third-party tools to build automation around scaling the application up and down, based on Stackdriver CPU usage monitoring.

Answer: B

Explanation:

A managed instance group can help use virtual machines directly and with autoscaling can scaling as per the demand.Refer GCP documentation. Managed Instance Groups AutoScaling Managed instance groups offer autoscaling capabilities that allow you to automatically add or delete instances from a managed instance group based on increases or decreases in load.

NO.11 You need to enable traffic between multiple groups of Compute Engine instances that are currently running two different GCP projects. Each group of Compute Engine instances is running in its own VPC. What should you do?

A. Verify that both projects are in a GCP Organization. Create a new VPC and add all instances.

B. Verify that both projects are in a GCP Organization. Share the VPC from one project and request that the Compute Engine instances in the other project use this shared VPC.

C. Verify that you are the Project Administrator of both projects. Create two new VPCs and add all instances.

D. Verify that you are the Project Administrator of both projects. Create a new VPC and add all instances.

Answer: B

Explanation:

For example, an existing instance in a service project cannot be reconfigured to use a Shared VPC network, but a new instance can be created to use available subnets in a Shared VPC network.
https://cloud.google.com/vpc/docs/shared-vpc

NO.12 You manage a VPC network in Google Cloud with a subnet that is rapidly approaching its private IP address capacity. You expect the number of Compute Engine VM instances in the same region to double within a week. You need to implement a Google-recommended solution that minimizes operational costs and does not require downtime. What should you do?

A. Create a second VPC with the same subnet IP range, and connect this VPC to the existing VPC by using VPC Network Peering.

B. Delete the existing subnet, and create a new subnet with double the IP range available.

C. Use the Google Cloud CLI tool to expand the primary IP range of your subnet.

D. Permit additional traffic from the expected range of private IP addresses to reach your VMs by configuring firewall rules.

Answer: C

Explanation:
Google Cloud allows you to expand the primary IP range of an existing subnet without downtime. This is the recommended and most cost-effective way to increase available private IP addresses while keeping all current resources online.

NO.13 You create a Deployment with 2 replicas in a Google Kubernetes Engine cluster that has a single preemptible node pool. After a few minutes, you use kubectlto examine the status of your Pod and observe that one of them is still in Pendingstatus:

```
$ kubectl get pods -l app=myapp
NAME                                READY     STATUS    RESTART   AGE
myapp-deployment-58ddbbb995-lp86m   0/1       Pending   0         9m
myapp-deployment-58ddbbb995-qjpkg   1/1       Running   0         9m
```

What is the most likely cause?

A. The pending Pod's resource requests are too large to fit on a single node of the cluster.

B. Too many Pods are already running in the cluster, and there are not enough resources left to schedule the pending Pod.

C. The node pool is configured with a service account that does not have permission to pull the container image used by the pending Pod.

D. The pending Pod was originally scheduled on a node that has been preempted between the creation of the Deployment and your verification of the Pods' status. It is currently being rescheduled on a new node.

Answer: B
Explanation:
Reasons for a Pod Status Pending:
Troubleshooting Reason #1: Not enough CPU
Troubleshooting Reason #2: Not enough memory
Troubleshooting Reason #3: Not enough CPU and memory
https://managedkube.com/kubernetes/k8sbot/troubleshooting/pending/pod/2019/02/22/pending-pod.html

NO.14 Your company is moving its entire workload to Compute Engine. Some servers should be accessible through the Internet, and other servers should only be accessible over the internal network. All servers need to be able to talk to each other over specific ports and protocols. The current on-premises network relies on a demilitarized zone (DMZ) for the public servers and a Local Area Network (LAN) for the private servers. You need to design the networking infrastructure on Google Cloud to match these requirements. What should you do?

A. 1. Create a single VPC with a subnet for the DMZ and a subnet for the LAN.
2. Set up firewall rules to open up relevant traffic between the DMZ and the LAN subnets, and another firewall rule to allow public ingress traffic for the DMZ.

B. 1. Create a single VPC with a subnet for the DMZ and a subnet for the LAN.
2. Set up firewall rules to open up relevant traffic between the DMZ and the LAN subnets, and another firewall rule to allow public egress traffic for the DMZ.

C. 1. Create a VPC with a subnet for the DMZ and another VPC with a subnet for the LAN.
2. Set up firewall rules to open up relevant traffic between the DMZ and the LAN subnets, and

another firewall rule to allow public ingress traffic for the DMZ.

D. 1. Create a VPC with a subnet for the DMZ and another VPC with a subnet for the LAN.
2. Set up firewall rules to open up relevant traffic between the DMZ and the LAN subnets, and another firewall rule to allow public egress traffic for the DMZ.

Answer: A

Explanation:
By default traffic between subnets on a VPC network is not allowed (except on the "default" network).
(This blocks traffic between all instances, not just traffic between subnets => FW rules must be defined to allow communications between all instances, regardless the subnets)
2 VPC will not work without peering.

NO.15 An application generates daily reports in a Compute Engine virtual machine (VM). The VM is in the project corp-iot-insights. Your team operates only in the project corp-aggregate-reports and needs a copy of the daily exports in the bucket corp-aggregate-reports-storage. You want to configure access so that the daily reports from the VM are available in the bucket corp-aggregate-reports-storage and use as few steps as possible while following Google-recommended practices. What should you do?

A. Move both projects under the same folder.

B. Grant the VM Service Account the role Storage Object Creator on corp-aggregate-reports- storage.

C. Create a Shared VPC network between both projects.
Grant the VM Service Account the role Storage Object Creator on corp-iot-insights.

D. Make corp-aggregate-reports-storage public and create a folder with a pseudo-randomized suffix name.
Share the folder with the IoT team.

Answer: B

Explanation:
Predefined roles
The following table describes Identity and Access Management (IAM) roles that are associated with Cloud Storage and lists the permissions that are contained in each role. Unless otherwise noted, these roles can be applied either to entire projects or specific buckets.
Storage Object Creator (roles/storage.objectCreator) Allows users to create objects. Does not give permission to view, delete, or overwrite objects.
https://cloud.google.com/storage/docs/access-control/iam-roles#standard-roles

NO.16 Your company has an internal application for managing transactional orders. The application is used exclusively by employees in a single physical location. The application requires strong consistency, fast queries, and ACID guarantees for multi-table transactional updates. The first version of the application is implemented inPostgreSQL, and you want to deploy it to the cloud with minimal code changes. Which database is most appropriate for this application?

A. BigQuery

B. Cloud SQL

C. Cloud Spanner

D. Cloud Datastore

Answer: B

Explanation:
Cloud SQL for PostgreSQL is a fully-managed database service that helps you set up, maintain, manage, and administer your PostgreSQL relational databases on Google Cloud Platform.
https://cloud.google.com/sql/docs/postgres

NO.17 You host your website on Compute Engine. The number of global users visiting your website is rapidly expanding. You need to minimize latency and support user growth in multiple geographical regions. You also want to follow Google-recommended practices and minimize operational costs. Which two actions should you take? Choose 2 answers

A. Deploy all of your VMs in a single Google Cloud region with the largest available CIDR range.

B. Deploy your VMs in multiple Google Cloud regions closest to your users' geographical locations.

C. Use an external Application Load Balancer in Regional mode.

D. Use an external Application Load Balancer in Global mode.

E. Use a Network Load Balancer.

Answer: BD

Explanation:
Placing VMs in multiple regions close to users reduces latency and scales better for a global audience. A Global external Application Load Balancer provides a single anycast IP and intelligently routes traffic to the closest healthy backend across regions, aligning with Google-recommended, cost-effective global load balancing practices.

NO.18 What is the gcloud command to create a cluster named ch09-cluster-10 with four nodes?

A. gcloud container clusters create ch09-cluster-10 4

B. gcloud container clusters create ch09-cluster-10 --num-nodes=4

C. gcloud clusters container create ch09-cluster-10 --num-nodes=4

D. gcloud container clusters create ch09-cluster-10 --nodes-num=4

Answer: B

NO.19 You are working in a team that has developed a new application that needs to be deployed on Kubernetes. The production application is business critical and should be optimized for reliability. You need to provision a Kubernetes cluster and want to follow Google-recommended practices. What should you do?

A. Create a GKE Autopilot cluster. Enroll the cluster in the rapid release channel.

B. Create a GKE Autopilot cluster. Enroll the cluster in the stable release channel.

C. Create a zonal GKE standard cluster. Enroll the cluster in the stable release channel.

D. Create a regional GKE standard cluster. Enroll the cluster in the rapid release channel.

Answer: B

Explanation:
"recommended practices" --> Autopilot
"optimized for reliability" --> Stable release

NO.20 You need to deploy an application in Google Cloud using serverless technology. You want to test a new version of the application with a small percentage of production traffic. What should you do?

A. Deploy the application to Cloud Run. Use gradual rollouts for traffic splitting.
B. Deploy the application to Google Kubernetes Engine. Use Anthos Service Mash for traffic splitting.
C. Deploy the application to Cloud Functions. Specify the version number in the functions name.
D. Deploy the application to App Engine. For each new version, create a new service.
Answer: A

NO.21 You want to select and configure a cost-effective solution for relational data on Google Cloud Platform. You are working with a small set of operational data in one geographic location. You need to support point- in-time recovery. What should you do?
A. Select Cloud SQL (MySQL). Verify that the enable binary logging option is selected.
B. Select Cloud SQL (MySQL). Select the create failover replicas option.
C. Select Cloud Spanner. Set up your instance with 2 nodes.
D. Select Cloud Spanner. Set up your instance as multi-regional.
Answer: A
Explanation:
You must enable binary logging to use point-in-time recovery. Enabling binary logging causes a slight reduction in write performance.
https://cloud.google.com/sql/docs/mysql/backup-recovery/backups

NO.22 Your company is seeking a scalable solution to retain and explore application logs hosted on Compute Engine. You must be able to analyze your logs with SQL queries, and you want to be able to create charts to identify patterns and trends in your logs over time. You want to follow Google-recommended practices and minimize your operational costs. What should you do?
A. Use a custom script to push your application logs to BigQuery for exploration.
B. Ingest your application logs to Cloud Logging by using Ops Agent, and explore your logs in Logs Explorer.
C. Ingest your application logs to Cloud Logging by using Ops Agent, and explore your logs with Log Analytics.
D. Use a custom script to push your application logs to Cloud SQL for exploration.
Answer: C
Explanation:
Ingesting your application logs to Cloud Logging using the Ops Agent and exploring the logs with Log Analytics is the recommended solution. This approach leverages Cloud Logging's built-in capabilities for querying logs with SQL-like syntax and allows you to visualize trends over time, all while minimizing operational costs. Log Analytics provides a powerful interface for log exploration and analysis.

NO.23 Your company was recently impacted by a service disruption that caused multiple Dataflow jobs to get stuck, resulting in significant downtime in downstream applications and revenue loss. You were able to resolve the issue by identifying and fixing an error you found in the code. You need to design a solution with minimal management effort to identify when jobs are stuck in the future to ensure that this issue does not occur again. What should you do?
A. Set up Error Reporting to identify stack traces that indicate slowdowns in Dataflow jobs. Set up alerts based on these log entries.

B. Use the Personalized Service Health dashboard to identify issues with Dataflow jobs across regions.
C. Update the Dataflow job configurations to send messages to a Pub/Sub topic when there are delays. Configure a backup Dataflow job to process jobs that are delayed. Use Cloud Tasks to trigger an alert when messages are pushed to the Pub/Sub topic.
D. Set up Cloud Monitoring alerts on the data freshness metric for the Dataflow jobs to receive a notification when a certain threshold is reached.
Answer: D
Explanation:
Setting up Cloud Monitoring alerts on the data freshness metric for Dataflow jobs allows you to monitor job performance and receive notifications when jobs fall behind or get stuck, based on a predefined threshold. This provides an efficient and automated way to detect delays without requiring complex manual intervention or additional jobs.

NO.24 You are given a project with a single virtual private cloud (VPC) and a single subnetwork in the us-central1 region. There is a Compute Engine instance hosting an application in this subnetwork. You need to deploy a new instance in the same project in the europe-west1 region.
This new instance needs access to the application. You want to follow Google-recommended practices. What should you do?
A. 1. Create a subnetwork in the same VPC, in europe-west1.
2. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.
B. 1. Create a VPC and a subnetwork in europe-west1.
2. Expose the application with an internal load balancer.
3. Create the new instance in the new subnetwork and use the load balancer's address as the endpoint.
C. 1. Create a subnetwork in the same VPC, in europe-west1.
2. Use Cloud VPN to connect the two subnetworks.
3. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.
D. 1. Create a VPC and a subnetwork in europe-west1.
2. Peer the 2 VPCs.
3. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.
Answer: A
Explanation:
Cloud VPN is used when connecting on-premise infrastrucutre to the cloud. In this case creating a new subnet in a new region, and using the internal IP to communicate is sufficient enough to fulfillthe requirements.

NO.25 You are using multiple configurations for gcloud. You want to review the configured Kubernetes Engine cluster of an inactive configuration using the fewest possible steps. What should you do?
A. Use gcloud config configurations describe to review the output.
B. Use gcloud config configurations activate and gcloud config list to review the output.

C. Use kubectl config get-contexts to review the output.

D. Use kubectl config use-context and kubectl config view to review the output.

Answer: D

Explanation:

Use kubectl config get-contexts to review the output: shows the clusters and the configurations and based on the output we can identify the inactive configurations.

https://kubernetes.io/docs/reference/kubectl/cheatsheet/#kubectl-context-and-configuration

NO.26 While working on a project, an application administrator has been given the responsibility of managing all resources. He wants to delegate the responsibility of managing the existing service accounts to another administrator. He will also be responsible to manage the other service accounts that will be created. Which of the following is the best way to delegate the privileges required to manage all the service accounts?

A. Granting iam.serviceAccountUser to the administrator at the project level

B. Granting iam.serviceProjectAccountUser to the administrator at the project level

C. Granting iam.serviceAccountUser to the administrator at the service account level

D. Granting iam.serviceProjectAccountUser to the administrator at the service account level

Answer: A

NO.27 You are building an archival solution for your data warehouse and have selected Cloud Storage to archive your data. Your users need to be able to access this archived data once a quarter for some regulatory requirements. You want to select a cost-efficient option. Which storage option should you use?

A. Cold Storage

B. Nearline Storage

C. Regional Storage

D. Multi-Regional Storage

Answer: B

Explanation:

Nearline, Coldline, and Archive offer ultra low-cost, highly-durable, highly available archival storage. For data accessed less than once a year, Archive is a cost-effective storage option for long-term preservation of data. Coldline is also ideal for cold storage--data your business expects to touch less than once a quarter. For warmer storage, choose Nearline: data you expect to access less than once a month, but possibly multiple times throughout the year. All storage classes are available across all GCP regions and provide unparalleled sub-second access speeds with a consistent API.

https://cloud.google.com/storage/archival

NO.28 You want to configure autohealing for network load balancing for a group of Compute Engine instances that run in multiple zones, using the fewest possible steps. You need to configure re-creation of VMs if they are unresponsive after 3 attempts of 10 seconds each. What should you do?

A. Create an HTTP load balancer with a backend configuration that references an existing instance group.

Set the health check to healthy (HTTP)

B. Create an HTTP load balancer with a backend configuration that references an existing instance

group.
Define a balancing mode and set the maximum RPS to 10.

C. Create a managed instance group. Set the Autohealing health check to healthy (HTTP)

D. Create a managed instance group. Verify that the autoscaling setting is on.

Answer: C

Explanation:
Use separate health checks for load balancing and for autohealing. Health checks for load balancing detect unresponsive instances and direct traffic away from them. Health checks for autohealing detect and recreate failed instances, so they should be less aggressive than load balancing health checks. Using the same health check for these services would remove the distinction between unresponsive instances and failed instances, causing unnecessary latency and unavailability for your users.
https://cloud.google.com/compute/docs/tutorials/high-availability-autohealing

NO.29 You have an application that uses Cloud Spanner as a backend database. The application has a very predictable traffic pattern. You want to automatically scale up or down the number of Spanner nodes depending on traffic. What should you do?

A. Create a cron job that runs on a scheduled basis to review stackdriver monitoring metrics, and then resize the Spanner instance accordingly.

B. Create a Stackdriver alerting policy to send an alert to oncall SRE emails when Cloud Spanner CPU exceeds the threshold. SREs would scale resources up or down accordingly.

C. Create a Stackdriver alerting policy to send an alert to Google Cloud Support email when Cloud Spanner CPU exceeds your threshold. Google support would scale resources up or down accordingly.

D. Create a Stackdriver alerting policy to send an alert to webhook when Cloud Spanner CPU is over or under your threshold. Create a Cloud Function that listens to HTTP and resizes Spanner resources accordingly.

Answer: D

Explanation:
Note: You can scale the number of nodes in your instance based on the Cloud Monitoring metrics on CPU or storage utilization in conjunction with Cloud Functions.
https://cloud.google.com/spanner/docs/instances

NO.30 You want to create a new role for your colleagues that will apply to all current and future projects created in your organization. The role should have the permissions of the BigQuery Job User and Cloud Bigtable User roles. You want to follow Google's recommended practices. How should you create the new role?

A. Use "gcloud iam combine-roles --global" to combine the 2 roles into a new custom role.

B. For one of your projects, in the Google Cloud Platform Console under Roles, select both roles and combine them into a new custom role. Use "gcloud iam promote-role" to promote the role from a project role to an organization role.

C. For all projects, in the Google Cloud Platform Console under Roles, select both roles and combine them into a new custom role.

D. For your organization, in the Google Cloud Platform Console under Roles, select both roles and combine them into a new custom role.

Answer: D
Explanation:
A is not correct because this does not create a new role.
B is not correct because gcloud cannot promote a role to org level.
C is not correct because it's recommended to define the role on the organization level. Also, the role will not be applied on new projects.
D is correct because this creates a new role with the combined permissions on the organization level.

NO.31 Your company has a rapidly growing social media platform and a user base primarily located in North America. Due to increasing demand, your current on-premises PostgreSQL database, hosted in your United States headquarters data center, no longer meets your needs. You need to identify a cloud-based database solution that offers automatic scaling, multi-region support for future expansion, and maintains low latency. What should you do?

A. Use Bigtable.

B. Use BigQuery.

C. Use Spanner.

D. Use Cloud SQL for PostgreSQL.

Answer: C
Explanation:
Using Cloud Spanner is the best option for a rapidly growing social media platform that requires automatic scaling, multi-region support, and low latency. Cloud Spanner is a fully managed, horizontally scalable relational database that supports strong consistency and multi-region replication, making it ideal for high-demand applications that require scalability and low-latency performance across regions. It also offers PostgreSQL compatibility for easy migration from your existing database.

NO.32 Your VMs are running in a subnet that has a subnet mask of 255.255.255.240. The current subnet has no more free IP addresses and you require an additional 10 IP addresses for new VMs. The existing and new VMs should all be able to reach each other without additional routes.
What should you do?

A. Use gcloud to expand the IP range of the current subnet.

B. Delete the subnet, and recreate it using a wider range of IP addresses.

C. Create a new project. Use Shared VPC to share the current network with the new project.

D. Create a new subnet with the same starting IP but a wider range to overwrite the current subnet.

Answer: A
Explanation:
https://cloud.google.com/sdk/gcloud/reference/compute/networks/subnets/expand-ip-range

NO.33 You are in charge of provisioning access for all Google Cloud users in your organization. Your company recently acquired a startup company that has their own Google Cloud organization. You need to ensure that your Site Reliability Engineers (SREs) have the same project permissions in the startup company's organization as in your own organization. What should you do?

A. In the Google Cloud console for your organization, select Create role from selection, and choose destination as the startup company's organization.

B. In the Google Cloud console for the startup company, select Create role from selection and choose

source as the startup company's Google Cloud organization.
C. Use the gcloud iam roles copy command, and provide the Organization ID of the startup company's Google Cloud Organization as the destination.
D. Use the gcloud iam roles copy command, and provide the project IDs of all projects in the startup company's organization as the destination.
Answer: C

NO.34 You are working for a hospital that stores its medical images in an on-premises data room. The hospital wants to use Cloud Storage for archival storage of these images. The hospital wants an automated process to upload any new medical images to Cloud Storage. You need to design and implement a solution. What should you do?
A. Create a Pub/Sub topic, and enable a Cloud Storage trigger for the Pub/Sub topic. Create an application that sends all medical images to the Pub/Sub topic.
B. Create a script that uses the gcloud storage command to synchronize the on-premises storage with Cloud Storage, Schedule the script as a cron job.
C. Create a Pub/Sub topic, and create a Cloud Function connected to the topic that writes data to Cloud Storage. Create an application that sends all medical images to the Pub/Sub topic.
D. In the Google Cloud console, go to Cloud Storage. Upload the relevant images to the appropriate bucket.
Answer: B

NO.35 You have been asked to migrate a docker application from datacenter to cloud. Your solution architect has suggested uploading docker images to GCR in one project and running an application in a GKE cluster in a separate project. You want to store images in the project img-278322 and run the application in the project prod-278986. You want to tag the image as acme_track_n_trace:v1. You want to follow Google-recommended practices. What should you do?
A. Run gcloud builds submit --tag gcr.io/img-278322/acme_track_n_trace
B. Run gcloud builds submit --tag gcr.io/img-278322/acme_track_n_trace:v1
C. Run gcloud builds submit --tag gcr.io/prod-278986/acme_track_n_trace
D. Run gcloud builds submit --tag gcr.io/prod-278986/acme_track_n_trace:v1
Answer: B
Explanation:
Run gcloud builds submit tag gcr.io/img-278322/acme_track_n_trace:v1. is the right answer. This command correctly tags the image as acme_track_n_trace:v1 and uploads the image to the img-278322 project.