[Jan 15, 2023 Genuine Associate-Cloud-Engineer Exam Dumps New 2023 Google Pratice Exam [Q102-Q116

[Jan 15, 2023] Genuine Associate-Cloud-Engineer Exam Dumps New 2023 Google Pratice Exam
New 2023 Realistic Associate-Cloud-Engineer Dumps Test Engine Exam Questions in here

NO.102 You are running an application on multiple virtual machines within a managed instance group and have autoscaling enabled. The autoscaling policy is configured so that additional instances are added to the group if the CPU utilization of instances goes above 80%. VMs are added until the instance group reaches its maximum limit of five VMs or until CPU utilization of instances lowers to 80%. The initial delay for HTTP health checks against the instances is set to 30 seconds. The virtual machine instances take around three minutes to become available for users. You observe that when the instance group autoscales, it adds more instances then necessary to support the levels of end-user traffic. You want to properly maintain instance group sizes when autoscaling.

What should you do?
* Set the maximum number of instances to 1.
* Decrease the maximum number of instances to 3.
* Use a TCP health check instead of an HTTP health check.
* Increase the initial delay of the HTTP health check to 200 seconds.

The virtual machine instances take around three minutes to become available for users.

NO.103 You need to create a copy of a custom Compute Engine virtual machine (VM) to facilitate an expected increase in application traffic due to a business acquisition. What should you do?
* Create a Compute Engine snapshot of your base VM. Create your images from that snapshot.
* Create a Compute Engine snapshot of your base VM. Create your instances from that snapshot.
* Create a custom Compute Engine image from a snapshot. Create your images from that image.
* Create a custom Compute Engine image from a snapshot. Create your instances from that image.
A custom image belongs only to your project. To create an instance with a custom image, you must first have a custom image.

NO.104 Your company has an existing GCP organization with hundreds of projects and a billing account.

Your company recently acquired another company that also has hundreds of projects and its own billing account. You would like to consolidate all GCP costs of both GCP organizations onto a single invoice. You would like to consolidate all costs as of tomorrow. What should you do?
* Link the acquired company’s projects to your company’s billing account.
* Configure the acquired company’s billing account and your company’s billing account to export the billing data into the same BigQuery dataset.
* Migrate the acquired company’s projects into your company’s GCP organization. Link the migrated projects to your company’s billing account.
* Create a new GCP organization and a new billing account. Migrate the acquired company’s projects and your company’s projects into the new GCP organization and link the projects to the new billing account.
Projects are linked to another organization as well in the acquired company so migrating would need google cloud support.

We can not do ourselves. however, we can link other company projects to an existing billing account to generate total cost.

https://cloud.google.com/resource-manager/docs/project-migration

NO.105 You’re trying to provide temporary access to some files in a Cloud Storage bucket. You want to limit the time that the files are available to 10 minutes. With the fewest steps possible, what is the best way to generate a signed URL?
* In the UI select the objects and click the Generate Signed URL button.
* Create a service account and JSON key. Use the gsutil signurl -t 10m command and pass in the JSON key and bucket.
* In the UI select the objects and click the “Sign With Key” button.
* Create a service account and JSON key. Use the gsutil signurl -d 10m command and pass in the JSON key and bucket.

NO.106 You are using Google Kubernetes Engine with autoscaling enabled to host a new application. You want to expose this new application to the public, using HTTPS on a public IP address. What should you do?
* Create a Kubernetes Service of type NodePort for your application, and a Kubernetes Ingress to expose this Service via a Cloud Load Balancer.
* Create a Kubernetes Service of type ClusterIP for your application. Configure the public DNS name of your application using the IP of this Service.
* Create a Kubernetes Service of type NodePort to expose the application on port 443 of each node of the Kubernetes cluster. Configure the public DNS name of your application with the IP of every node of the cluster to achieve load-balancing.
* Create a HAProxy pod in the cluster to load-balance the traffic to all the pods of the application. Forward the public traffic to HAProxy with an iptable rule. Configure the DNS name of your application using the public IP of the node HAProxy is running on.
Reference:

https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer

NO.107 A Solutions Architect is designing a high-performance computing job that runs on Amazon EC2 instances in private subnets. To allow the application to download patches, the infrastructure must be altered to allow the instances to access external

endpoints. Any changes to the infrastructure must involve minimal ongoing systems management effort.

What will allow the EC2 instances to access the endpoint while meeting these requirements?
* NAT gateway
* Elastic IP address
* AWS Direct Connect
* Virtual private gateway
Explanation/Reference: https://aws.amazon.com/vpc/

NO.108 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?
* Set up a policy that uses Nearline storage for 30 days and then moves to Archive storage for three years.
* Set up a policy that uses Standard storage for 30 days and then moves to Archive storage for three years.
* 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.
* 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.
Reference:

https://books.google.com.pk/books?id=q0nhDwAAQBAJ&pg=PA52&lpg=PA52&dq=Set+up+a

+policy+that+uses+Nearline+storage+for+30+days+and+then+moves+to+Archive+storage+for+three

+years.&source=bl&ots=kYLZN1ymA8&sig=ACfU3U2XLmzQ39cmPDwjfWxRbNtDNLc_6g&hl=en&sa=X&ved

=2ahUKEwjZmefOpr7qAhVzQkEAHTUgASYQ6AEwAHoECAoQAQ#v=onepage&q=Set%20up%20a%
20policy%20that%20uses%20Nearline%20storage%20for%2030%20days%20and%20then%20moves%20to

%20Archive%20storage%20for%20three%20years.&f=false

NO.109 Users submit requests to a service that takes several minutes to process. A Solutions Architect needs to ensure that these requests are processed at least once, and that the service has the ability to handle large increases in the number of requests.

How should these requirements be met?
* Put the requests into an Amazon SQS queue and configure Amazon EC2 instances to poll the queue
* Publish the message to an Amazon SNS topic that an Amazon EC2 subscriber can receive and process
* Save the requests to an Amazon DynamoDB table with a DynamoDB stream that triggers an Amazon EC2 Spot Instance
* Use Amazon S3 to store the requests and configure an event notification to have Amazon EC2 instances process the new object

NO.110 You’ve setup and tested several custom roles in your development project. What is the fastest way to create the same roles for your new production project?
* Recreate them in the new project.
* Use the gcloud iam copy roles command and set the destination project.
* In the UI, select the roles and click the Export button.
* Use the gcloud iam roles copy command and set the destination project.

NO.111 Your security team has been reluctant to move to the cloud because they don’t have the level of network visibility they’re used to. Which feature might help them to gain insights into your Google Cloud network?
* Routes

* Subnets
* Flow logs
* Firewall rules

NO.112 You are hosting an application on bare-metal servers in your own data center. The application needs access to Cloud Storage. However, security policies prevent the servers hosting the application from having public IP addresses or access to the internet. You want to follow Google-recommended practices to provide the application with access to Cloud Storage. What should you do?
* 1. Use nslookup to get the IP address for storage.googleapis.com.2. Negotiate with the security team to be able to give a public IP address to the servers.3. Only allow egress traffic from those servers to the IP addresses for storage.googleapis.com.
* 1. Using Cloud VPN, create a VPN tunnel to a Virtual Private Cloud (VPC) in Google Cloud Platform (GCP).2. In this VPC, create a Compute Engine instance and install the Squid proxy server on this instance.3. Configure your servers to use that instance as a proxy to access Cloud Storage.
* 1. Use Migrate for Compute Engine (formerly known as Velostrata) to migrate those servers to Compute Engine.2. Create an internal load balancer (ILB) that uses storage.googleapis.com as backend.3. Configure your new instances to use this ILB as proxy.
* 1. Using Cloud VPN or Interconnect, create a tunnel to a VPC in GCP.2. Use Cloud Router to create a custom route advertisement for 199.36.153.4/30. Announce that network to your on-premises network through the VPN tunnel.3. In your on-premises network, configure your DNS server to resolve *.googleapis.com as a CNAME to restricted.googleapis.com.

NO.113 You are the organization and billing administrator for your company. The engineering team has the Project Creator role on the organization. You do not want the engineering team to be able to link projects to the billing account. Only the finance team should be able to link a project to a billing account, but they should not be able to make any other changes to projects. What should you do?
* Assign the finance team only the Billing Account User role on the billing account.
* Assign the engineering team only the Billing Account User role on the billing account.
* Assign the finance team the Billing Account User role on the billing account and the Project Billing Manager role on the organization.
* Assign the engineering team the Billing Account User role on the billing account and the Project Billing Manager role on the organization.

NO.114 You have an application that receives SSL-encrypted TCP traffic on port 443. Clients for this application are located all over the world. You want to minimize latency for the clients. Which load balancing option should you use?
* HTTPS Load Balancer
* Network Load Balancer
* SSL Proxy Load Balancer
* Internal TCP/UDP Load Balancer. Add a firewall rule allowing ingress traffic from 0.0.0.0/0on the target instances.
Explanation/Reference: https://cloud.google.com/load-balancing/docs/ssl

NO.115 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?
* Select Cloud SQL (MySQL). Verify that the enable binary logging option is selected.
* Select Cloud SQL (MySQL). Select the create failover replicas option.
* Select Cloud Spanner. Set up your instance with 2 nodes.
* Select Cloud Spanner. Set up your instance as multi-regional.

NO.116 Users of your application are complaining of slowness when loading the application. You realize the slowness is because the App Engine deployment serving the application is deployed in us-central whereas all users of this application are closest to europe-west3. You want to change the region of the App Engine application to europe-west3 to minimize latency. What’s the best way to change the App Engine region?
* Create a new project and create an App Engine instance in europe-west3

* Use the gcloud app region set command and supply the name of the new region.
* From the console, under the App Engine page, click edit, and change the region drop-down.
* Contact Google Cloud Support and request the change.

App engine is a regional service, which means the infrastructure that runs your app(s) is located in a specific region and is managed by Google to be redundantly available across all the zones within that region. Once an app engine deployment is created in a region, it cant be changed. The only way is to create a new project and create an App Engine instance in europe-west3, send all user traffic to this instance and delete the app engine instance in us-central.

Ref: https://cloud.google.com/appengine/docs/locations

Target audience and other prerequisites

This certification is intended for the Cloud Engineers who know how to plan and configure Cloud solutions, install Cloud solution environments, implement and deploy Cloud solutions, configure security and access, ensure the successful operation of Cloud solutions.

There are no official requirements for this certificate and its qualifying test. However, before taking the certification exam, it is recommended that you possess at least six months of practical experience with Google Cloud solutions and products.

Deploying & Implementing Cloud Solutions - Deploy and implement Google Kubernetes Engine Resources: this domain includes skills in deploying Google Kubernetes Engine cluster, container application to Google Kubernetes Engine with pods, and configuring Google Kubernetes Engine Application logging and monitoring;- Deploy and implement data solutions: the applicants should be able to initialize data systems with Google products and load data;- Deploy an application infrastructure with Cloud Deployment Manager with a focus on the development and launch of Deployment Manager templates.- Deploy and implement compute engine resources: the skills required for this area include launching compute instances with Cloud SDK and Cloud Console; creating auto-scaled managed instance groups with instance templates; uploading/generating custom SSH keys for instances; configuring virtual machines or Stackdriver logging and monitoring. It also requires that the students can assess compute quotas and request for increases; install a Stackdriver Agent for logging and monitoring;- Deploy and implement networking resources: this will measure your ability to launch Compute Engine instances using custom network configuration; create egress and ingress firewall rule or VPC; create VPNs between the external network and Google VPC with Cloud VPN; create the load balancer to allocate the application network traffic for apps; **Grab latest Amazon Associate-Cloud-Engineer Dumps as PDF Updated:** https://www.actualtests4sure.com/Associate-Cloud-Engineer-test-questions.html]