Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer

NEW QUESTION 1
Your managed instance group raised an alert stating that new instance creation has failed to create new instances. You need to maintain the number of running instances specified by the template to be able to process expected application traffic. What should you do?

A. Create an instance template that contains valid syntax which will be used by the instance grou
B. Delete any persistent disks with the same name as instance names.
C. Create an instance template that contains valid syntax that will be used by the instance grou
D. Verify that the instance name and persistent disk name values are not the same in the template.
E. Verify that the instance template being used by the instance group contains valid synta
F. Delete any persistent disks with the same name as instance name
G. Set the disks.autoDelete property to true in the instance template.
H. Delete the current instance template and replace it with a new instance templat
I. Verify that the instance name and persistent disk name values are not the same in the templat
J. Set the disks.autoDelete property to true in the instance template.

Answer: C

NEW QUESTION 2
You are running multiple VPC-native Google Kubernetes Engine clusters in the same subnet. The IPs available for the nodes are exhausted, and you want to ensure that the clusters can grow in nodes when needed. What should you do?

A. Create a new subnet in the same region as the subnet being used.
B. Add an alias IP range to the subnet used by the GKE clusters.
C. Create a new VPC, and set up VPC peering with the existing VPC.
D. Expand the CIDR range of the relevant subnet for the cluster.

Answer: C

Explanation:
To create a VPC peering connection, first create a request to peer with another VPC.

NEW QUESTION 3
You are building a new version of an application hosted in an App Engine environment. You want to test the new version with 1% of users before you completely switch your application over to the new version. What should you do?

A. Deploy a new version of your application in Google Kubernetes Engine instead of App Engine and then use GCP Console to split traffic.
B. Deploy a new version of your application in a Compute Engine instance instead of App Engine and then use GCP Console to split traffic.
C. Deploy a new version as a separate app in App Engin
D. Then configure App Engine using GCP Console to split traffic between the two apps.
E. Deploy a new version of your application in App Engin
F. Then go to App Engine settings in GCP Console and split traffic between the current version and newly deployed versions accordingly.

Answer: A

NEW QUESTION 4
You need to immediately change the storage class of an existing Google Cloud bucket. You need to reduce service cost for infrequently accessed files stored in that bucket and for all files that will be added to that bucket in the future. What should you do?

A. Use the gsutil to rewrite the storage class for the bucket Change the default storage class for the bucket
B. Use the gsutil to rewrite the storage class for the bucket Set up Object Lifecycle management on the bucket
C. Create a new bucket and change the default storage class for the bucket Set up Object Lifecycle management on lite bucket
D. Create a new bucket and change the default storage class for the bucket import the files from the previous bucket into the new bucket

Answer: B

NEW QUESTION 5
You have a development project with appropriate IAM roles defined. You are creating a production project and want to have the same IAM roles on the new project, using the fewest possible steps. What should you do?

A. Use gcloud iam roles copy and specify the production project as the destination project.
B. Use gcloud iam roles copy and specify your organization as the destination organization.
C. In the Google Cloud Platform Console, use the 'create role from role' functionality.
D. In the Google Cloud Platform Console, use the 'create role' functionality and select all applicable permissions.

Answer: B

NEW QUESTION 6
Your projects incurred more costs than you expected last month. Your research reveals that a development GKE container emitted a huge number of logs, which resulted in higher costs. You want to disable the logs quickly using the minimum number of steps. What should you do?

A. 1. Go to the Logs ingestion window in Stackdriver Logging, and disable the log source for the GKE container resource.
B. 1. Go to the Logs ingestion window in Stackdriver Logging, and disable the log source for the GKE Cluster Operations resource.
C. 1. Go to the GKE console, and delete existing clusters.2. Recreate a new cluster.3. Clear the option to enable legacy Stackdriver Logging.
D. 1. Go to the GKE console, and delete existing clusters.2. Recreate a new cluster.3. Clear the option to enable legacy Stackdriver Monitoring.

Answer: A

NEW QUESTION 7
Your company publishes large files on an Apache web server that runs on a Compute Engine instance. The Apache web server is not the only application running in the project. You want to receive an email when the egress network costs for the server exceed 100 dollars for the current month as measured by Google Cloud Platform (GCP). What should you do?

A. Set up a budget alert on the project with an amount of 100 dollars, a threshold of 100%, and notificationtype of "email."
B. Set up a budget alert on the billing account with an amount of 100 dollars, a threshold of 100%, and notification type of "email."
C. Export the billing data to BigQuer
D. Create a Cloud Function that uses BigQuery to sum the egress network costs of the exported billing data for the Apache web server for the current month and sends an email if it is over 100 dollar
E. Schedule the Cloud Function using Cloud Scheduler to run hourly.
F. Use the Stackdriver Logging Agent to export the Apache web server logs to Stackdriver Logging.Create a Cloud Function that uses BigQuery to parse the HTTP response log data in Stackdriver for the current month and sends an email if the size of all HTTP responses, multiplied by current GCP egress prices, totals over 100 dollar
G. Schedule the Cloud Function using Cloud Scheduler to run hourly.

Answer: D

NEW QUESTION 8
You have an instance group that you want to load balance. You want the load balancer to terminate the client SSL session. The instance group is used to serve a public web application over HTTPS. You want to follow Google-recommended practices. What should you do?

A. Configure an HTTP(S) load balancer.
B. Configure an internal TCP load balancer.
C. Configure an external SSL proxy load balancer.
D. Configure an external TCP proxy load balancer.

Answer: A

NEW QUESTION 9
You have deployed an application on a single Compute Engine instance. The application writes logs to disk. Users start reporting errors with the application. You want to diagnose the problem. What should you do?

A. Navigate to Cloud Logging and view the application logs.
B. Connect to the instance's serial console and read the application logs.
C. Configure a Health Check on the instance and set a Low Healthy Threshold value.
D. Install and configure the Cloud Logging Agent and view the logs from Cloud Logging.

Answer: D

NEW QUESTION 10
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
C. SREs would scale resources up or down accordingly.
D. Create a Stackdriver alerting policy to send an alert to Google Cloud Support email when Cloud Spanner CPU exceeds your threshol
E. Google support would scale resources up or down accordingly.
F. Create a Stackdriver alerting policy to send an alert to webhook when Cloud Spanner CPU is over or under your threshold
G. Create a Cloud Function that listens to HTTP and resizes Spanner resources accordingly.

Answer: D

NEW QUESTION 10
Your company is moving from an on-premises environment to Google Cloud Platform (GCP). You have
multiple development teams that use Cassandra environments as backend databases. They all need a development environment that is isolated from other Cassandra instances. You want to move to GCP quickly and with minimal support effort. What should you do?

A. * 1. Build an instruction guide to install Cassandra on GCP.* 2. Make the instruction guide accessible to your developers.
B. * 1. Advise your developers to go to Cloud Marketplace.* 2. Ask the developers to launch a Cassandra image for their development work.
C. * 1. Build a Cassandra Compute Engine instance and take a snapshot of it.* 2. Use the snapshot to create instances for your developers.
D. * 1. Build a Cassandra Compute Engine instance and take a snapshot of it.* 2. Upload the snapshot to Cloud Storage and make it accessible to your developers.* 3. Build instructions to create a Compute Engine instance from the snapshot so that developers can do it themselves.

Answer: D

NEW QUESTION 13
You have successfully created a development environment in a project for an application. This application uses Compute Engine and Cloud SQL. Now, you need to create a production environment for this application.
The security team has forbidden the existence of network routes between these 2 environments, and asks you to follow Google-recommended practices. What should you do?

A. Create a new project, enable the Compute Engine and Cloud SQL APIs in that project, and replicate the setup you have created in the development environment.
B. Create a new production subnet in the existing VPC and a new production Cloud SQL instance in your existing project, and deploy your application using those resources.

C. Create a new project, modify your existing VPC to be a Shared VPC, share that VPC with your new project, and replicate the setup you have in the development environment in that new project, in the Shared VPC.
D. Ask the security team to grant you the Project Editor role in an existing production project used by another division of your compan
E. Once they grant you that role, replicate the setup you have in the development environment in that project.

Answer: A

NEW QUESTION 16
You will have several applications running on different Compute Engine instances in the same project. You want to specify at a more granular level the service account each instance uses when calling Google Cloud APIs. What should you do?

A. When creating the instances, specify a Service Account for each instance
B. When creating the instances, assign the name of each Service Account as instance metadata
C. After starting the instances, use gcloud compute instances update to specify a Service Account for each instance
D. After starting the instances, use gcloud compute instances update to assign the name of the relevant Service Account as instance metadata

Answer: C

NEW QUESTION 17
You have deployed multiple Linux instances on Compute Engine. You plan on adding more instances in the coming weeks. You want to be able to access all of these instances through your SSH client over me Internet without having to configure specific access on the existing and new instances. You do not want the Compute Engine instances to have a public IP. What should you do?

A. Configure Cloud Identity-Aware Proxy (or HTTPS resources
B. Configure Cloud Identity-Aware Proxy for SSH and TCP resources.
C. Create an SSH keypair and store the public key as a project-wide SSH Key
D. Create an SSH keypair and store the private key as a project-wide SSH Key

Answer: C

NEW QUESTION 18
You need a dynamic way of provisioning VMs on Compute Engine. The exact specifications will be in a dedicated configuration file. You want to follow Google's recommended practices. Which method should you use?

A. Deployment Manager
B. Cloud Composer
C. Managed Instance Group
D. Unmanaged Instance Group

Answer: C

NEW QUESTION 19
You need to provide a cost estimate for a Kubernetes cluster using the GCP pricing calculator for Kubernetes. Your workload requires high IOPs, and you will also be using disk snapshots. You start by entering the number of nodes, average hours, and average days. What should you do next?

A. Fill in local SS
B. Fill in persistent disk storage and snapshot storage.
C. Fill in local SS
D. Add estimated cost for cluster management.
E. Select Add GPU
F. Fill in persistent disk storage and snapshot storage.
G. Select Add GPU
H. Add estimated cost for cluster management.

Answer: C

NEW QUESTION 22
You are setting up a Windows VM on Compute Engine and want to make sure you can log in to the VM via RDP. What should you do?

A. After the VM has been created, use your Google Account credentials to log in into the VM.
B. After the VM has been created, use gcloud compute reset-windows-password to retrieve the login credentials for the VM.
C. When creating the VM, add metadata to the instance using 'windows-password' as the key and a password as the value.
D. After the VM has been created, download the JSON private key for the default Compute Engine service accoun
E. Use the credentials in the JSON file to log in to the VM.

Answer: D

NEW QUESTION 25
You built an application on Google Cloud Platform that uses Cloud Spanner. Your support team needs to monitor the environment but should not have access to table data. You need a streamlined solution to grant the correct permissions to your support team, and you want to follow Google-recommended practices. What should you do?

A. Add the support team group to the roles/monitoring.viewer role
B. Add the support team group to the roles/spanner.databaseUser role.
C. Add the support team group to the roles/spanner.databaseReader role.
D. Add the support team group to the roles/stackdriver.accounts.viewer role.

Answer: B

NEW QUESTION 29
You have a Linux VM that must connect to Cloud SQL. You created a service account with the appropriate access rights. You want to make sure that the VM uses this service account instead of the default Compute Engine service account. What should you do?

A. When creating the VM via the web console, specify the service account under the 'Identity and API Access' section.
B. Download a JSON Private Key for the service accoun
C. On the Project Metadata, add that JSON as the value for the key compute-engine-service-account.
D. Download a JSON Private Key for the service accoun
E. On the Custom Metadata of the VM, add that JSON as the value for the key compute-engine-service-account.
F. Download a JSON Private Key for the service accoun
G. After creating the VM, ssh into the VM and save the JSON under ~/.gcloud/compute-engine-service-account.json.

Answer: A

NEW QUESTION 34
You deployed an LDAP server on Compute Engine that is reachable via TLS through port 636 using UDP. You want to make sure it is reachable by clients over that port. What should you do?

A. Add the network tag allow-udp-636 to the VM instance running the LDAP server.
B. Create a route called allow-udp-636 and set the next hop to be the VM instance running the LDAP server.
C. Add a network tag of your choice to the instanc
D. Create a firewall rule to allow ingress on UDP port 636 for that network tag.
E. Add a network tag of your choice to the instance running the LDAP serve
F. Create a firewall rule to allow egress on UDP port 636 for that network tag.

Answer: C

NEW QUESTION 39
You need to assign a Cloud Identity and Access Management (Cloud IAM) role to an external auditor. The auditor needs to have permissions to review your Google Cloud Platform (GCP) Audit Logs and also to review your Data Access logs. What should you do?

A. Assign the auditor the IAM role roles/logging.privateLogViewe
B. Perform the export of logs to Cloud Storage.
C. Assign the auditor the IAM role roles/logging.privateLogViewe
D. Direct the auditor to also review the logs for changes to Cloud IAM policy.
E. Assign the auditor's IAM user to a custom role that has logging.privateLogEntries.list permissio
F. Perform the export of logs to Cloud Storage.
G. Assign the auditor's IAM user to a custom role that has logging.privateLogEntries.list permissio
H. Direct the auditor to also review the logs for changes to Cloud IAM policy.

Answer: C

NEW QUESTION 42
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 kubectl to examine the status of your Pod and observe that one of them is still in Pending status:

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' statu
E. It is currently being rescheduled on a new node.

Answer: B

NEW QUESTION 44
You have been asked to set up the billing configuration for a new Google Cloud customer. Your customer wants to group resources that share common IAM policies. What should you do?

A. Use labels to group resources that share common IAM policies
B. Use folders to group resources that share common IAM policies
C. Set up a proper billing account structure to group IAM policies
D. Set up a proper project naming structure to group IAM policies

Answer: B

NEW QUESTION 46
You want to run a single caching HTTP reverse proxy on GCP for a latency-sensitive website. This specific reverse proxy consumes almost no CPU. You want to have a 30-GB in-memory cache, and need an additional 2 GB of memory for the rest of the processes. You want to minimize cost. How should you run this reverse proxy?

A. Create a Cloud Memorystore for Redis instance with 32-GB capacity.
B. Run it on Compute Engine, and choose a custom instance type with 6 vCPUs and 32 GB of memory.
C. Package it in a container image, and run it on Kubernetes Engine, using n1-standard-32 instances as nodes.

D. Run it on Compute Engine, choose the instance type n1-standard-1, and add an SSD persistent disk of 32 GB.

Answer: B

NEW QUESTION 47
You need to deploy an application, which is packaged in a container image, in a new project. The application exposes an HTTP endpoint and receives very few requests per day. You want to minimize costs. What should you do?

A. Deploy the container on Cloud Run.
B. Deploy the container on Cloud Run on GKE.
C. Deploy the container on App Engine Flexible.
D. Deploy the container on Google Kubernetes Engine, with cluster autoscaling and horizontal pod autoscaling enabled.

Answer: A

NEW QUESTION 50
Your organization has three existing Google Cloud projects. You need to bill the Marketing department for only their Google Cloud services for a new initiative within their group. What should you do?

A. * 1. Verify that you ace assigned the Billing Administrator IAM role tor your organization's Google Cloud Project for the Marketing department* 2. Link the new project to a Marketing Billing Account
B. * 1. Verify that you are assigned the Billing Administrator IAM role for your organization's Google Cloud account* 2. Create a new Google Cloud Project for the Marketing department* 3. Set the default key-value project labels to department marketing for all services in this project
C. * 1. Verify that you are assigned the Organization Administrator IAM role for your organization's Google Cloud account* 2. Create a new Google Cloud Project for the Marketing department 3. Link the new project to a Marketing Billing Account.
D. * 1. Verity that you are assigned the Organization Administrator IAM role for your organization's Google Cloud account* 2. Create a new Google Cloud Project for the Marketing department* 3. Set the default key value project labels to department marketing for all services in this protect

Answer: A

NEW QUESTION 52
You received a JSON file that contained a private key of a Service Account in order to get access to several resources in a Google Cloud project. You downloaded and installed the Cloud SDK and want to use this private key for authentication and authorization when performing gcloud commands. What should you do?

A. Use the command gcloud auth login and point it to the private key
B. Use the command gcloud auth activate-service-account and point it to the private key
C. Place the private key file in the installation directory of the Cloud SDK and rename it to "credentials ison"
D. Place the private key file in your home directory and rename it to ''GOOGLE_APPUCATION_CREDENTiALS".

Answer: A

NEW QUESTION 56
Your company has a 3-tier solution running on Compute Engine. The configuration of the current infrastructure is shown below.

Each tier has a service account that is associated with all instances within it. You need to enable communication on TCP port 8080 between tiers as follows:
• Instances in tier #1 must communicate with tier #2.
• Instances in tier #2 must communicate with tier #3. What should you do?

A. 1. Create an ingress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.2.0/24)• Protocols: allow all2. Create an ingress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.1.0/24)•Protocols: allow all
B. 1. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #2 service account• Source filter: all instances with tier #1 service account• Protocols: allow TCP:80802. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #3 service account• Source filter: all instances with tier #2 service account• Protocols: allow TCP: 8080
C. 1. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #2 service account• Source filter: all instances with tier #1 service account• Protocols: allow all2. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #3 service account• Source filter: all instances with tier #2 service account• Protocols: allow all
D. 1. Create an egress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.2.0/24)• Protocols: allow TCP: 80802. Create an egress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.1.0/24)• Protocols: allow TCP: 8080

Answer: B

NEW QUESTION 61
Your organization is a financial company that needs to store audit log files for 3 years. Your organization has hundreds of Google Cloud projects. You need to implement a cost-effective approach for log file retention. What should you do?

A. Create an export to the sink that saves logs from Cloud Audit to BigQuery.
B. Create an export to the sink that saves logs from Cloud Audit to a Coldline Storage bucket.
C. Write a custom script that uses logging API to copy the logs from Stackdriver logs to BigQuery.
D. Export these logs to Cloud Pub/Sub and write a Cloud Dataflow pipeline to store logs to Cloud SQL.

Answer: A

NEW QUESTION 64
Your auditor wants to view your organization's use of data in Google Cloud. The auditor is most interested in auditing who accessed data in Cloud Storage buckets. You need to help the auditor access the data they need. What should you do?

A. Assign the appropriate permissions, and then use Cloud Monitoring to review metrics
B. Use the export logs API to provide the Admin Activity Audit Logs in the format they want
C. Turn on Data Access Logs for the buckets they want to audit, and Then build a query in the log viewer that filters on Cloud Storage
D. Assign the appropriate permissions, and then create a Data Studio report on Admin Activity Audit Logs

Answer: A

NEW QUESTION 67
You want to verify the IAM users and roles assigned within a GCP project named my-project. What should you do?

A. Run gcloud iam roles lis
B. Review the output section.
C. Run gcloud iam service-accounts lis
D. Review the output section.
E. Navigate to the project and then to the IAM section in the GCP Consol
F. Review the members and roles.
G. Navigate to the project and then to the Roles section in the GCP Consol
H. Review the roles and status.

Answer: D

NEW QUESTION 71
Your company uses a large number of Google Cloud services centralized in a single project. All teams have specific projects for testing and development. The DevOps team needs access to all of the production services in order to perform their job. You want to prevent Google Cloud product changes from broadening their permissions in the future. You want to follow Google-recommended practices. What should you do?

A. Grant all members of the DevOps team the role of Project Editor on the organization level.
B. Grant all members of the DevOps team the role of Project Editor on the production project.
C. Create a custom role that combines the required permission
D. Grant the DevOps team the custom role on the production project.
E. Create a custom role that combines the required permission
F. Grant the DevOps team the custom role on the organization level.

Answer: A

NEW QUESTION 74
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 instance
B. Set the 'Automatic Restart' to o
C. Set the 'On-host maintenance' to Migrate VM instanc
D. Add the instance template to an instance group.
E. Create an instance template for the instance
F. Set 'Automatic Restart' to of
G. Set 'On-host maintenance' to Terminate VM instance
H. Add the instance template to an instance group.
I. Create an instance group for the instance
J. Set the 'Autohealing' health check to healthy (HTTP).
K. Create an instance group for the instanc
L. Verify that the 'Advanced creation options' setting for 'do not retry machine creation' is set to off.

Answer: B

NEW QUESTION 75
You need to create a new billing account and then link it with an existing Google Cloud Platform project. What should you do?

A. Verify that you are Project Billing Manager for the GCP projec
B. Update the existing project to link it to the existing billing account.
C. Verify that you are Project Billing Manager for the GCP projec
D. Create a new billing account and linkthe new billing account to the existing project.
E. Verify that you are Billing Administrator for the billing accoun
F. Create a new project and link the new project to the existing billing account.
G. Verify that you are Billing Administrator for the billing accoun
H. Update the existing project to link it to the existing billing account.

Answer: C

NEW QUESTION 80
Your organization needs to grant users access to query datasets in BigQuery but prevent them from accidentally deleting the datasets. You want a solution that follows Google-recommended practices. What should you do?

A. Add users to roles/bigquery user role only, instead of roles/bigquery dataOwner.
B. Add users to roles/bigquery dataEditor role only, instead of roles/bigquery dataOwner.
C. Create a custom role by removing delete permissions, and add users to that role only.
D. Create a custom role by removing delete permission
E. Add users to the group, and then add the group to the custom role.

Answer: B

NEW QUESTION 84
Your existing application running in Google Kubernetes Engine (GKE) consists of multiple pods running on four GKE n1–standard–2 nodes. You need to deploy additional pods requiring n2–highmem–16 nodes without any downtime. What should you do?

A. Use gcloud container clusters upgrad
B. Deploy the new services.
C. Create a new Node Pool and specify machine type n2–highmem–16. Deploy the new pods.
D. Create a new cluster with n2–highmem–16 node
E. Redeploy the pods and delete the old cluster.
F. Create a new cluster with both n1–standard–2 and n2–highmem–16 node
G. Redeploy the pods and delete the old cluster.

Answer: B

NEW QUESTION 89
You need to manage multiple Google Cloud Platform (GCP) projects in the fewest steps possible. You want to configure the Google Cloud SDK command line interface (CLI) so that you can easily manage multiple GCP projects. What should you?

A. * 1. Create a configuration for each project you need to manage.* 2. Activate the appropriate configuration when you work with each of your assigned GCP projects.
B. * 1. Create a configuration for each project you need to manage.* 2. Use gcloud init to update the configuration values when you need to work with a non-default project
C. * 1. Use the default configuration for one project you need to manage.* 2. Activate the appropriate configuration when you work with each of your assigned GCP projects.
D. * 1. Use the default configuration for one project you need to manage.* 2. Use gcloud init to update the configuration values when you need to work with a non-default project.

Answer: D

NEW QUESTION 91
You have a Compute Engine instance hosting a production application. You want to receive an email if the instance consumes more than 90% of its CPU resources for more than 15 minutes. You want to use Google services. What should you do?

A. * 1. Create a consumer Gmail account.* 2. Write a script that monitors the CPU usage.* 3. When the CPU usage exceeds the threshold, have that script send an email using the Gmail account and smtp.gmail.com on port 25 as SMTP server.
B. * 1. Create a Stackdriver Workspace, and associate your Google Cloud Platform (GCP) project with it.* 2. Create an Alerting Policy in Stackdriver that uses the threshold as a trigger conditio
C. * 3. Configure your email address in the notification channel.
D. * 1. Create a Stackdriver Workspace, and associate your GCP project with it.* 2. Write a script that monitors the CPU usage and sends it as a custom metric to Stackdrive
E. * 3. Create an uptime check for the instance in Stackdriver.
F. * 1. In Stackdriver Logging, create a logs-based metric to extract the CPU usage by using this regular expression: CPU Usage: ([0-9] {1,3})%* 2. In Stackdriver Monitoring, create an Alerting Policy based on this metri
G. * 3. Configure your email address in the notification channel.

Answer: D

NEW QUESTION 94
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 projec
D. Use Shared VPC to share the current network with the new project.
E. Create a new subnet with the same starting IP but a wider range to overwrite the current subnet.

Answer: C

NEW QUESTION 96
You have an application on a general-purpose Compute Engine instance that is experiencing excessive disk read throttling on its Zonal SSD Persistent Disk. The application primarily reads large files from disk. The disk size is currently 350 GB. You want to provide the maximum amount of throughput while minimizing costs. What should you do?

A. Increase the size of the disk to 1 TB.
B. Increase the allocated CPU to the instance.
C. Migrate to use a Local SSD on the instance.
D. Migrate to use a Regional SSD on the instance.

Answer: C

NEW QUESTION 97
You have one project called proj-sa where you manage all your service accounts. You want to be able to use a service account from this project to take snapshots of VMs running in another project called proj-vm. What should you do?

A. Download the private key from the service account, and add it to each VMs custom metadata.
B. Download the private key from the service account, and add the private key to each VM's SSH keys.
C. Grant the service account the IAM Role of Compute Storage Admin in the project called proj-vm.
D. When creating the VMs, set the service account's API scope for Compute Engine to read/write.

Answer: C

NEW QUESTION 102
You are deploying an application to a Compute Engine VM in a managed instance group. The application must be running at all times, but only a single instance of the VM should run per GCP project. How should you configure the instance group?

A. Set autoscaling to On, set the minimum number of instances to 1, and then set the maximum number of instances to 1.
B. Set autoscaling to Off, set the minimum number of instances to 1, and then set the maximum number of instances to 1.
C. Set autoscaling to On, set the minimum number of instances to 1, and then set the maximum number of instances to 2.
D. Set autoscaling to Off, set the minimum number of instances to 1, and then set the maximum number of instances to 2.

Answer: A

NEW QUESTION 106
You need to set a budget alert for use of Compute Engineer services on one of the three Google Cloud Platform projects that you manage. All three projects are linked to a single billing account. What should you do?

A. Verify that you are the project billing administrato
B. Select the associated billing account and create a budget and alert for the appropriate project.
C. Verify that you are the project billing administrato
D. Select the associated billing account and create a budget and a custom alert.
E. Verify that you are the project administrato
F. Select the associated billing account and create a budget for the appropriate project.
G. Verify that you are project administrato
H. Select the associated billing account and create a budget and a custom alert.

Answer: B

NEW QUESTION 111
You are creating an application that will run on Google Kubernetes Engine. You have identified MongoDB as the most suitable database system for your application and want to deploy a managed MongoDB environment that provides a support SLA. What should you do?

A. Create a Cloud Bigtable cluster and use the HBase API
B. Deploy MongoDB Alias from the Google Cloud Marketplace
C. Download a MongoDB installation package and run it on Compute Engine instances
D. Download a MongoDB installation package, and run it on a Managed Instance Group

Answer: D

NEW QUESTION 112
You have designed a solution on Google Cloud Platform (GCP) that uses multiple GCP products. Your company has asked you to estimate the costs of the solution. You need to provide estimates for the monthly total cost. What should you do?

A. For each GCP product in the solution, review the pricing details on the products pricing pag
B. Use the pricing calculator to total the monthly costs for each GCP product.
C. For each GCP product in the solution, review the pricing details on the products pricing pag
D. Create a Google Sheet that summarizes the expected monthly costs for each product.
E. Provision the solution on GC
F. Leave the solution provisioned for 1 wee
G. Navigate to the Billing Report page in the Google Cloud Platform Consol
H. Multiply the 1 week cost to determine the monthly costs.
I. Provision the solution on GC
J. Leave the solution provisioned for 1 wee
K. Use Stackdriver to determine the provisioned and used resource amount
L. Multiply the 1 week cost to determine the monthly costs.

Answer: A

NEW QUESTION 115
You deployed an App Engine application using gcloud app deploy, but it did not deploy to the intended project. You want to find out why this happened and where the application deployed. What should you do?

A. Check the app.yaml file for your application and check project settings.
B. Check the web-application.xml file for your application and check project settings.
C. Go to Deployment Manager and review settings for deployment of applications.
D. Go to Cloud Shell and run gcloud config list to review the Google Cloud configuration used for deployment.

Answer: A

NEW QUESTION 120
You need to add a group of new users to Cloud Identity. Some of the users already have existing Google accounts. You want to follow one of Google's recommended practices and avoid conflicting accounts. What should you do?

A. Invite the user to transfer their existing account
B. Invite the user to use an email alias to resolve the conflict
C. Tell the user that they must delete their existing account
D. Tell the user to remove all personal email from the existing account

Answer: B

NEW QUESTION 121
An employee was terminated, but their access to Google Cloud Platform (GCP) was not removed until 2 weeks later. You need to find out this employee accessed any sensitive customer information after their termination. What should you do?

A. View System Event Logs in Stackdrive
B. Search for the user's email as the principal.
C. View System Event Logs in Stackdrive
D. Search for the service account associated with the user.
E. View Data Access audit logs in Stackdrive
F. Search for the user's email as the principal.
G. View the Admin Activity log in Stackdrive
H. Search for the service account associated with the user.

Answer: B

NEW QUESTION 124
You have created an application that is packaged into a Docker image. You want to deploy the Docker image as a workload on Google Kubernetes Engine. What should you do?

A. Upload the image to Cloud Storage and create a Kubernetes Service referencing the image
B. Upload the image to Cloud Storage and create a Kubernetes Deployment referencing the image
C. Upload the image to Container Registry and create a Kubernetes Service referencing the image.
D. Upload the image to Container Registry and create a Kubernetes Deployment referencing the mage

Answer: C

NEW QUESTION 128
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.

NEW QUESTION 133
You are configuring service accounts for an application that spans multiple projects. Virtual machines (VMs) running in the web-applications project need access to BigQuery datasets in crm-databases-proj. You want to follow Google-recommended practices to give access to the service account in the web-applications project. What should you do?

A. Give "project owner" for web-applications appropriate roles to crm-databases- proj
B. Give "project owner" role to crm-databases-proj and the web-applications project.
C. Give "project owner" role to crm-databases-proj and bigquery.dataViewer role to web-applications.
D. Give bigquery.dataViewer role to crm-databases-proj and appropriate roles to web-applications.

Answer: D

NEW QUESTION 136
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?

A. 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.
B. 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.
C. 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.
D. 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.

Answer: D

NEW QUESTION 141
You want to add a new auditor to a Google Cloud Platform project. The auditor should be allowed to read, but not modify, all project items.
How should you configure the auditor's permissions?

A. Create a custom role with view-only project permission
B. Add the user's account to the custom role.
C. Create a custom role with view-only service permission
D. Add the user's account to the custom role.
E. Select the built-in IAM project Viewer rol
F. Add the user's account to this role.
G. Select the built-in IAM service Viewer rol
H. Add the user's account to this role.

Answer: C

NEW QUESTION 145
You created several resources in multiple Google Cloud projects. All projects are linked to different billing accounts. To better estimate future charges, you want to have a single visual representation of all costs incurred. You want to include new cost data as soon as possible. What should you do?

A. Configure Billing Data Export to BigQuery and visualize the data in Data Studio.
B. Visit the Cost Table page to get a CSV export and visualize it using Data Studio.
C. Fill all resources in the Pricing Calculator to get an estimate of the monthly cost.
D. Use the Reports view in the Cloud Billing Console to view the desired cost information.

Answer: A

NEW QUESTION 147
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 Spanne
D. Set up your instance with 2 nodes.
E. Select Cloud Spanne
F. Set up your instance as multi-regional.

Answer: A

NEW QUESTION 152
You have sensitive data stored in three Cloud Storage buckets and have enabled data access logging. You want to verify activities for a particular user for these buckets, using the fewest possible steps. You need to verify the addition of metadata labels and which files have been viewed from those buckets. What should you do?

A. Using the GCP Console, filter the Activity log to view the information.
B. Using the GCP Console, filter the Stackdriver log to view the information.
C. View the bucket in the Storage section of the GCP Console.
D. Create a trace in Stackdriver to view the information.

Answer: A

NEW QUESTION 154
You are building an application that will run in your data center. The application will use Google Cloud Platform (GCP) services like AutoML. You created a service account that has appropriate access to AutoML. You need to enable authentication to the APIs from your on-premises environment. What should you do?

A. Use service account credentials in your on-premises application.
B. Use gcloud to create a key file for the service account that has appropriate permissions.
C. Set up direct interconnect between your data center and Google Cloud Platform to enable authentication for your on-premises applications.
D. Go to the IAM & admin console, grant a user account permissions similar to the service account permissions, and use this user account for authentication from your data center.

Answer: B

NEW QUESTION 157
You have a developer laptop with the Cloud SDK installed on Ubuntu. The Cloud SDK was installed from the Google Cloud Ubuntu package repository. You want to test your application locally on your laptop with Cloud Datastore. What should you do?

A. Export Cloud Datastore data using gcloud datastore export.
B. Create a Cloud Datastore index using gcloud datastore indexes create.
C. Install the google-cloud-sdk-datastore-emulator component using the apt get install command.
D. Install the cloud-datastore-emulator component using the gcloud components install command.

Answer: D

NEW QUESTION 159
You need to produce a list of the enabled Google Cloud Platform APIs for a GCP project using the gcloud command line in the Cloud Shell. The project name is my-project. What should you do?

A. Run gcloud projects list to get the project ID, and then run gcloud services list --project <project ID>.
B. Run gcloud init to set the current project to my-project, and then run gcloud services list --available.
C. Run gcloud info to view the account value, and then run gcloud services list --account <Account>.
D. Run gcloud projects describe <project ID> to verify the project value, and then run gcloud services list--available.

Answer: A

NEW QUESTION 160
Your organization uses Active Directory (AD) to manage user identities. Each user uses this identity for federated access to various on-premises systems. Your security team has adopted a policy that requires users to log into Google Cloud with their AD identity instead of their own login. You want to follow the Google-recommended practices to implement this policy. What should you do?

A. Sync Identities with Cloud Directory Sync, and then enable SAML for single sign-on
B. Sync Identities in the Google Admin console, and then enable Oauth for single sign-on
C. Sync identities with 3rd party LDAP sync, and then copy passwords to allow simplified login with (he same credentials
D. Sync identities with Cloud Directory Sync, and then copy passwords to allow simplified login with the same credentials.

Answer: A

NEW QUESTION 161
You are building an application that stores relational data from users. Users across the globe will use this application. Your CTO is concerned about the scaling requirements because the size of the user base is unknown. You need to implement a database solution that can scale with your user growth with minimum configuration changes. Which storage solution should you use?

A. Cloud SQL
B. Cloud Spanner
C. Cloud Firestore
D. Cloud Datastore

Answer: D

NEW QUESTION 165
Your management has asked an external auditor to review all the resources in a specific project. The security team has enabled the Organization Policy called Domain Restricted Sharing on the organization node by specifying only your Cloud Identity domain. You want the auditor to only be able to view, but not modify, the resources in that project. What should you do?

A. Ask the auditor for their Google account, and give them the Viewer role on the project.
B. Ask the auditor for their Google account, and give them the Security Reviewer role on the project.
C. Create a temporary account for the auditor in Cloud Identity, and give that account the Viewer role on the project.
D. Create a temporary account for the auditor in Cloud Identity, and give that account the Security Reviewer role on the project.

Answer: A

NEW QUESTION 169
You are building an application that processes data files uploaded from thousands of suppliers. Your primary goals for the application are data security and the expiration of aged data. You need to design the application to:
•Restrict access so that suppliers can access only their own data.
•Give suppliers write access to data only for 30 minutes.
•Delete data that is over 45 days old.
You have a very short development cycle, and you need to make sure that the application requires minimal maintenance. Which two strategies should you use? (Choose two.)

A. Build a lifecycle policy to delete Cloud Storage objects after 45 days.
B. Use signed URLs to allow suppliers limited time access to store their objects.
C. Set up an SFTP server for your application, and create a separate user for each supplier.
D. Build a Cloud function that triggers a timer of 45 days to delete objects that have expired.
E. Develop a script that loops through all Cloud Storage buckets and deletes any buckets that are older than 45 days.

Answer: AE

NEW QUESTION 170
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: A

NEW QUESTION 175

Your company implemented BigQuery as an enterprise data warehouse. Users from multiple business units run queries on this data warehouse. However, you notice that query costs for BigQuery are very high, and you need to control costs. Which two methods should you use? (Choose two.)

A. Split the users from business units to multiple projects.
B. Apply a user- or project-level custom query quota for BigQuery data warehouse.
C. Create separate copies of your BigQuery data warehouse for each business unit.
D. Split your BigQuery data warehouse into multiple data warehouses for each business unit.
E. Change your BigQuery query model from on-demand to flat rat
F. Apply the appropriate number of slots to each Project.

Answer: BE

NEW QUESTION 176
You are running a data warehouse on BigQuery. A partner company is offering a recommendation engine based on the data in your data warehouse. The partner company is also running their application on Google Cloud. They manage the resources in their own project, but they need access to the BigQuery dataset in your project. You want to provide the partner company with access to the dataset What should you do?

A. Create a Service Account in your own project, and grant this Service Account access to BigGuery in your project
B. Create a Service Account in your own project, and ask the partner to grant this Service Account access to BigQuery in their project
C. Ask the partner to create a Service Account in their project, and have them give the Service Account access to BigQuery in their project
D. Ask the partner to create a Service Account in their project, and grant their Service Account access to the BigQuery dataset in your project

Answer: D

NEW QUESTION 180
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?

A. HTTPS Load Balancer
B. Network Load Balancer
C. SSL Proxy Load Balancer
D. Internal TCP/UDP Load Balance
E. Add a firewall rule allowing ingress traffic from 0.0.0.0/0 on the target instances.

Answer: C

NEW QUESTION 182
You are deploying a production application on Compute Engine. You want to prevent anyone from accidentally destroying the instance by clicking the wrong button. What should you do?

A. Disable the flag "Delete boot disk when instance is deleted."
B. Enable delete protection on the instance.
C. Disable Automatic restart on the instance.
D. Enable Preemptibility on the instance.

Answer: A

NEW QUESTION 186
Your company set up a complex organizational structure on Google Could Platform. The structure includes hundreds of folders and projects. Only a few team members should be able to view the hierarchical structure. You need to assign minimum permissions to these team members and you want to follow Google-recommended practices. What should you do?

A. Add the users to roles/browser role.
B. Add the users to roles/iam.roleViewer role.
C. Add the users to a group, and add this group to roles/browser role.
D. Add the users to a group, and add this group to roles/iam.roleViewer role.

Answer: C

NEW QUESTION 191
Your company has a large quantity of unstructured data in different file formats. You want to perform ETL transformations on the data. You need to make the data accessible on Google Cloud so it can be processed by a Dataflow job. What should you do?

A. Upload the data to BigQuery using the bq command line tool.
B. Upload the data to Cloud Storage using the gsutil command line tool.
C. Upload the data into Cloud SQL using the import function in the console.
D. Upload the data into Cloud Spanner using the import function in the console.

Answer: A

NEW QUESTION 194
For analysis purposes, you need to send all the logs from all of your Compute Engine instances to a BigQuery dataset called platform-logs. You have already installed the Stackdriver Logging agent on all the instances. You want to minimize cost. What should you do?

A. 1. Give the BigQuery Data Editor role on the platform-logs dataset to the service accounts used by your instances.2. Update your instances' metadata to add the following value: logs-destination:bq://platform-logs.
B. 1. In Stackdriver Logging, create a logs export with a Cloud Pub/Sub topic called logs as a sink.2.Create a Cloud Function that is triggered by messages in the

logs topic.3. Configure that Cloud Function to drop logs that are not from Compute Engine and to insert Compute Engine logs in the platform-logs dataset.
C. 1. In Stackdriver Logging, create a filter to view only Compute Engine logs.2. Click Create Export.3.Choose BigQuery as Sink Service, and the platform-logs dataset as Sink Destination.
D. 1. Create a Cloud Function that has the BigQuery User role on the platform-logs dataset.2. Configure this Cloud Function to create a BigQuery Job that executes this query:INSERT INTOdataset.platform-logs (timestamp, log)SELECT timestamp, log FROM compute.logsWHERE timestamp> DATE_SUB(CURRENT_DATE(), INTERVAL 1 DAY)3. Use Cloud Scheduler to trigger this Cloud Function once a day.

Answer: C

NEW QUESTION 196
You significantly changed a complex Deployment Manager template and want to confirm that the dependencies of all defined resources are properly met before committing it to the project. You want the most rapid feedback on your changes. What should you do?

A. Use granular logging statements within a Deployment Manager template authored in Python.
B. Monitor activity of the Deployment Manager execution on the Stackdriver Logging page of the GCP Console.
C. Execute the Deployment Manager template against a separate project with the same configuration, and monitor for failures.
D. Execute the Deployment Manager template using the –-preview option in the same project, and observe the state of interdependent resources.

Answer: D

NEW QUESTION 201
You need to set up a policy so that videos stored in a specific Cloud Storage Regional bucket are moved to Coldline after 90 days, and then deleted after one year from their creation. How should you set up the policy?

A. Use Cloud Storage Object Lifecycle Management using Age conditions with SetStorageClass andDelete action
B. Set the SetStorageClass action to 90 days and the Delete action to 275 days (365 – 90)
C. Use Cloud Storage Object Lifecycle Management using Age conditions with SetStorageClass and Delete action
D. Set the SetStorageClass action to 90 days and the Delete action to 365 days.
E. Use gsutil rewrite and set the Delete action to 275 days (365-90).
F. Use gsutil rewrite and set the Delete action to 365 days.

Answer: A

NEW QUESTION 205
You need to configure optimal data storage for files stored in Cloud Storage for minimal cost. The files are used in a mission-critical analytics pipeline that is used continually. The users are in Boston, MA (United States). What should you do?

A. Configure regional storage for the region closest to the users Configure a Nearline storage class
B. Configure regional storage for the region closest to the users Configure a Standard storage class
C. Configure dual-regional storage for the dual region closest to the users Configure a Nearline storage class
D. Configure dual-regional storage for the dual region closest to the users Configure a Standard storage class

Answer: B

NEW QUESTION 208
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

NEW QUESTION 213
You need to create an autoscaling managed instance group for an HTTPS web application. You want to make sure that unhealthy VMs are recreated. What should you do?

A. Create a health check on port 443 and use that when creating the Managed Instance Group.
B. Select Multi-Zone instead of Single-Zone when creating the Managed Instance Group.
C. In the Instance Template, add the label 'health-check'.
D. In the Instance Template, add a startup script that sends a heartbeat to the metadata server.

Answer: C

NEW QUESTION 215
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?

A. Set the maximum number of instances to 1.
B. Decrease the maximum number of instances to 3.
C. Use a TCP health check instead of an HTTP health check.

D. Increase the initial delay of the HTTP health check to 200 seconds.

Answer: D

NEW QUESTION 216
You are creating a Google Kubernetes Engine (GKE) cluster with a cluster autoscaler feature enabled. You need to make sure that each node of the cluster will run a monitoring pod that sends container metrics to a third-party monitoring solution. What should you do?

A. Deploy the monitoring pod in a StatefulSet object.
B. Deploy the monitoring pod in a DaemonSet object.
C. Reference the monitoring pod in a Deployment object.
D. Reference the monitoring pod in a cluster initializer at the GKE cluster creation time.

Answer: B

NEW QUESTION 221
......

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

Associate-Cloud-Engineer Practice Exam Features:

* Associate-Cloud-Engineer Questions and Answers Updated Frequently

* Associate-Cloud-Engineer Practice Questions Verified by Expert Senior Certified Staff

* Associate-Cloud-Engineer Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* Associate-Cloud-Engineer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The Associate-Cloud-Engineer Practice Test Here