Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer

NEW QUESTION 1
Your company runs one batch process in an on-premises server that takes around 30 hours to complete. The task runs monthly, can be performed offline, and must be restarted if interrupted. You want to migrate this workload to the cloud while minimizing cost. What should you do?

A. Migrate the workload to a Compute Engine Preemptible VM.
B. Migrate the workload to a Google Kubernetes Engine cluster with Preemptible nodes.
C. Migrate the workload to a Compute Engine V
D. Start and stop the instance as needed.
E. Create an Instance Template with Preemptible VMs O
F. Create a Managed Instance Group from the template and adjust Target CPU Utilizatio
G. Migrate the workload.

Answer: B

NEW QUESTION 2
You want to configure a solution for archiving data in a Cloud Storage bucket. The solution must be
cost-effective. Data with multiple versions should be archived after 30 days. Previous versions are accessed once a month for reporting. This archive data is also occasionally updated at month-end. What should you do?

A. Add a bucket lifecycle rule that archives data with newer versions after 30 days to Coldline Storage.
B. Add a bucket lifecycle rule that archives data with newer versions after 30 days to Nearline Storage.
C. Add a bucket lifecycle rule that archives data from regional storage after 30 days to Coldline Storage.
D. Add a bucket lifecycle rule that archives data from regional storage after 30 days to Nearline Storage.

Answer: B

NEW QUESTION 3
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 4
You manage an App Engine Service that aggregates and visualizes data from BigQuery. The application is deployed with the default App Engine Service account. The data that needs to be visualized resides in a different project managed by another team. You do not have access to this project, but you want your application to be able to read data from the BigQuery dataset. What should you do?

A. Ask the other team to grant your default App Engine Service account the role of BigQuery Job User.
B. Ask the other team to grant your default App Engine Service account the role of BigQuery Data Viewer.
C. In Cloud IAM of your project, ensure that the default App Engine service account has the role of BigQuery Data Viewer.
D. In Cloud IAM of your project, grant a newly created service account from the other team the role of BigQuery Job User in your project.

Answer: C

Explanation:
The Owner, Editor, and Viewer primitive roles include the BigQuery Admin (roles/bigquery.dataOwner), BigQuery Data Editor (roles/bigquery.dataEditor), and BigQuery Data Viewer (roles/bigquery.dataViewer) roles, respectively. This means the Owner, Editor, and Viewer primitive roles have BigQuery access as defined for the respective BigQuery roles.

NEW QUESTION 5
You are analyzing Google Cloud Platform service costs from three separate projects. You want to use this information to create service cost estimates by service type, daily and monthly, for the next six months using standard query syntax. What should you do?

A. Export your bill to a Cloud Storage bucket, and then import into Cloud Bigtable for analysis.
B. Export your bill to a Cloud Storage bucket, and then import into Google Sheets for analysis.
C. Export your transactions to a local file, and perform analysis with a desktop tool.
D. Export your bill to a BigQuery dataset, and then write time window-based SQL queries for analysis.

Answer: D

NEW QUESTION 6
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?

A. Assign the finance team only the Billing Account User role on the billing account.
B. Assign the engineering team only the Billing Account User role on the billing account.

C. Assign the finance team the Billing Account User role on the billing account and the Project Billing Manager role on the organization.
D. Assign the engineering team the Billing Account User role on the billing account and the Project Billing Manager role on the organization.

Answer: D

NEW QUESTION 7
Your company uses Cloud Storage to store application backup files for disaster recovery purposes. You want to follow Google's recommended practices. Which storage option should you use?

A. Multi-Regional Storage
B. Regional Storage
C. Nearline Storage
D. Coldline Storage

Answer: D

NEW QUESTION 8
You are hosting an application from Compute Engine virtual machines (VMs) in us–central1–a. You want to adjust your design to support the failure of a single Compute Engine zone, eliminate downtime, and minimize cost. What should you do?

A. – Create Compute Engine resources in us–central1–b.–Balance the load across both us–central1–a and us–central1–b.
B. – Create a Managed Instance Group and specify us–central1–a as the zone.–Configure the Health Check with a short Health Interval.
C. – Create an HTTP(S) Load Balancer.–Create one or more global forwarding rules to direct traffic to your VMs.
D. – Perform regular backups of your application.–Create a Cloud Monitoring Alert and be notified if your application becomes unavailable.–Restore from backups when notified.

Answer: C

NEW QUESTION 9
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 10
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 10
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 14
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 15
You need to set up permissions for a set of Compute Engine instances to enable them to write data into a particular Cloud Storage bucket. You want to follow Google-recommended practices. What should you do?

A. Create a service account with an access scop
B. Use the access scope 'https://www.googleapis.com/auth/devstorage.write_only'.
C. Create a service account with an access scop
D. Use the access scope 'https://www.googleapis.com/auth/cloud-platform'.
E. Create a service account and add it to the IAM role 'storage.objectCreator' for that bucket.
F. Create a service account and add it to the IAM role 'storage.objectAdmin' for that bucket.

Answer: B

NEW QUESTION 16
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 projec
D. Deploy the new version in that application.Use the App Engine library to proxy 1% of the requests to the new version.
E. Create a new App Engine application in the same projec
F. Deploy the new version in that application.Configure your network load balancer to send 1% of the traffic to that new application.

Answer: B

NEW QUESTION 19
You are asked to set up application performance monitoring on Google Cloud projects A, B, and C as a single pane of glass. You want to monitor CPU, memory, and disk. What should you do?

A. Enable API and then share charts from project A, B, and C.
B. Enable API and then give the metrics.reader role to projects A, B, and C.
C. Enable API and then use default dashboards to view all projects in sequence.
D. Enable API, create a workspace under project A, and then add project B and C.

Answer: D

NEW QUESTION 23
You need to run an important query in BigQuery but expect it to return a lot of records. You want to find out how much it will cost to run the query. You are using on-demand pricing. What should you do?

A. Arrange to switch to Flat-Rate pricing for this query, then move back to on-demand.
B. Use the command line to run a dry run query to estimate the number of bytes rea
C. Then convert that bytes estimate to dollars using the Pricing Calculator.
D. Use the command line to run a dry run query to estimate the number of bytes returne
E. Then convert that bytes estimate to dollars using the Pricing Calculator.
F. Run a select count (*) to get an idea of how many records your query will look throug
G. Then convert that number of rows to dollars using the Pricing Calculator.

Answer: B

NEW QUESTION 27
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 29
You have a batch workload that runs every night and uses a large number of virtual machines (VMs). It is fault- tolerant and can tolerate some of the VMs being terminated. The current cost of VMs is too high. What should you do?

A. Run a test using simulated maintenance event
B. If the test is successful, use preemptible N1 Standard VMs when running future jobs.
C. Run a test using simulated maintenance event

D. If the test is successful, use N1 Standard VMs when running future jobs.
E. Run a test using a managed instance grou
F. If the test is successful, use N1 Standard VMs in the managed instance group when running future jobs.
G. Run a test using N1 standard VMs instead of N2. If the test is successful, use N1 Standard VMs when running future jobs.

Answer: B

NEW QUESTION 32
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 36
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 38
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 39
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 44
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 48
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:

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

Answer: B

NEW QUESTION 53
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 54
You have a Compute Engine instance hosting an application used between 9 AM and 6 PM on weekdays. You want to back up this instance daily for disaster recovery purposes. You want to keep the backups for 30 days. You want the Google-recommended solution with the least management overhead and the least number of services. What should you do?

A. * 1. Update your instances' metadata to add the following value: snapshot–schedule: 0 1 * * ** 2. Update your instances' metadata to add the following value: snapshot–retention: 30
B. * 1. In the Cloud Console, go to the Compute Engine Disks page and select your instance's disk.* 2. In the Snapshot Schedule section, select Create Schedule and configure the following parameters:–Schedule frequency: Daily–Start time: 1:00 AM – 2:00 AM–Autodelete snapshots after 30 days
C. * 1. Create a Cloud Function that creates a snapshot of your instance's disk.* 2. Create a Cloud Function that deletes snapshots that are older than 30 day
D. 3.Use Cloud Scheduler to trigger both Cloud Functions daily at 1:00 AM.
E. * 1. Create a bash script in the instance that copies the content of the disk to Cloud Storage.* 2. Create a bash script in the instance that deletes data older than 30 days in the backup Cloud Storage bucket.* 3. Configure the instance's crontab to execute these scripts daily at 1:00 AM.

Answer: B

NEW QUESTION 59
You manage three Google Cloud projects with the Cloud Monitoring API enabled. You want to follow Google-recommended practices to visualize CPU and network metrics for all three projects together. What should you do?

A. * 1. Create a Cloud Monitoring Dashboard* 2. Collect metrics and publish them into the Pub/Sub topics 3. Add CPU and network Charts (or each of (he three projects
B. * 1. Create a Cloud Monitoring Dashboard.* 2. Select the CPU and Network metrics from the three projects.* 3. Add CPU and network Charts lot each of the three protects.
C. * 1 Create a Service Account and apply roles/viewer on the three projects* 2. Collect metrics and publish them lo the Cloud Monitoring API* 3. Add CPU and network Charts for each of the three projects.
D. * 1. Create a fourth Google Cloud project* 2 Create a Cloud Workspace from the fourth project and add the other three projects

Answer: B

NEW QUESTION 60
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 63
You want to deploy an application on Cloud Run that processes messages from a Cloud Pub/Sub topic. You want to follow Google-recommended practices. What should you do?

A. 1. Create a Cloud Function that uses a Cloud Pub/Sub trigger on that topic.2. Call your application on Cloud Run from the Cloud Function for every message.
B. 1. Grant the Pub/Sub Subscriber role to the service account used by Cloud Run.2. Create a Cloud Pub/Sub subscription for that topic.3. Make your application pull messages from that subscription.
C. 1. Create a service account.2. Give the Cloud Run Invoker role to that service account for your Cloud Run application.3. Create a Cloud Pub/Sub subscription

that uses that service account and uses your Cloud Run application as the push endpoint.
D. 1. Deploy your application on Cloud Run on GKE with the connectivity set to Internal.2. Create a Cloud Pub/Sub subscription for that topic.3. In the same Google Kubernetes Engine cluster as your application, deploy a container that takes the messages and sends them to your application.

Answer: D

NEW QUESTION 66
You need to select and configure compute resources for a set of batch processing jobs. These jobs take around 2 hours to complete and are run nightly. You want to minimize service costs. What should you do?

A. Select Google Kubernetes Engin
B. Use a single-node cluster with a small instance type.
C. Select Google Kubernetes Engin
D. Use a three-node cluster with micro instance types.
E. Select Compute Engin
F. Use preemptible VM instances of the appropriate standard machine type.
G. Select Compute Engin
H. Use VM instance types that support micro bursting.

Answer: C

NEW QUESTION 68
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?

A. Create a Kubernetes Service of type NodePort for your application, and a Kubernetes Ingress to expose this Service via a Cloud Load Balancer.
B. Create a Kubernetes Service of type ClusterIP for your applicatio
C. Configure the public DNS name of your application using the IP of this Service.
D. Create a Kubernetes Service of type NodePort to expose the application on port 443 of each node of the Kubernetes cluste
E. Configure the public DNS name of your application with the IP of every node of the cluster to achieve load-balancing.
F. 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 rul
G. Configure the DNS name of your application using the public IP of the node HAProxy is running on.

Answer: A

NEW QUESTION 73
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 78
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 79
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 82
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 85
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 88
You are managing a Data Warehouse on BigQuery. An external auditor will review your company's processes, and multiple external consultants will need view access to the data. You need to provide them with view access while following Google-recommended practices. What should you do?

A. Grant each individual external consultant the role of BigQuery Editor
B. Grant each individual external consultant the role of BigQuery Viewer
C. Create a Google Group that contains the consultants and grant the group the role of BigQuery Editor
D. Create a Google Group that contains the consultants, and grant the group the role of BigQuery Viewer

Answer: D

NEW QUESTION 92
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 94
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 98
You are storing sensitive information in a Cloud Storage bucket. For legal reasons, you need to be able to record all requests that read any of the stored data. You want to make sure you comply with these requirements. What should you do?

A. Enable the Identity Aware Proxy API on the project.
B. Scan the bucker using the Data Loss Prevention API.
C. Allow only a single Service Account access to read the data.
D. Enable Data Access audit logs for the Cloud Storage API.

Answer: D

NEW QUESTION 103
You need to configure IAM access audit logging in BigQuery for external auditors. You want to follow Google-recommended practices. What should you do?

A. Add the auditors group to the 'logging.viewer' and 'bigQuery.dataViewer' predefined IAM roles.
B. Add the auditors group to two new custom IAM roles.
C. Add the auditor user accounts to the 'logging.viewer' and 'bigQuery.dataViewer' predefined IAM roles.
D. Add the auditor user accounts to two new custom IAM roles.

Answer: C

NEW QUESTION 104
You are assigned to maintain a Google Kubernetes Engine (GKE) cluster named dev that was deployed on Google Cloud. You want to manage the GKE configuration using the command line interface (CLI). You have just downloaded and installed the Cloud SDK. You want to ensure that future CLI commands by default address this specific cluster. What should you do?

A. Use the command gcloud config sot container/cluster dev
B. Use the command gcloud container clusters update dev
C. Create a file called gk
D. default in the -/ .gcloud folder that contains the cluster name
E. Create a file called default
F. j son in the -/.gcioud folder that contains the cluster name

Answer: B

NEW QUESTION 107
You have 32 GB of data in a single file that you need to upload to a Nearline Storage bucket. The WAN connection you are using is rated at 1 Gbps, and you are the only one on the connection. You want to use as much of the rated 1 Gbps as possible to transfer the file rapidly. How should you upload the file?

A. Use the GCP Console to transfer the file instead of gsutil.
B. Enable parallel composite uploads using gsutil on the file transfer.
C. Decrease the TCP window size on the machine initiating the transfer.
D. Change the storage class of the bucket from Nearline to Multi-Regional.

Answer: B

NEW QUESTION 109
You want to configure an SSH connection to a single Compute Engine instance for users in the dev1 group. This instance is the only resource in this particular Google Cloud Platform project that the dev1 users should be able to connect to. What should you do?

A. Set metadata to enable-oslogin=true for the instanc
B. Grant the dev1 group the compute.osLogin role.Direct them to use the Cloud Shell to ssh to that instance.
C. Set metadata to enable-oslogin=true for the instanc
D. Set the service account to no service account for that instanc
E. Direct them to use the Cloud Shell to ssh to that instance.
F. Enable block project wide keys for the instanc
G. Generate an SSH key for each user in the dev1 group.Distribute the keys to dev1 users and direct them to use their third-party tools to connect.
H. Enable block project wide keys for the instanc
I. Generate an SSH key and associate the key with that instanc
J. Distribute the key to dev1 users and direct them to use their third-party tools to connect.

Answer: D

NEW QUESTION 114
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 116
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 118
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 120
You created an instance of SQL Server 2017 on Compute Engine to test features in the new version. You want to connect to this instance using the fewest number of steps. What should you do?

A. Install a RDP client on your deskto
B. Verify that a firewall rule for port 3389 exists.
C. Install a RDP client in your deskto
D. Set a Windows username and password in the GCP Consol
E. Use the credentials to log in to the instance.
F. Set a Windows password in the GCP Consol
G. Verify that a firewall rule for port 22 exist
H. Click theRDP button in the GCP Console and supply the credentials to log in.
I. Set a Windows username and password in the GCP Consol
J. Verify that a firewall rule for port 3389 exist
K. Click the RDP button in the GCP Console, and supply the credentials to log in.

Answer: B

NEW QUESTION 125
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 128
Your organization has user identities in Active Directory. Your organization wants to use Active Directory as their source of truth for identities. Your organization wants to have full control over the Google accounts used by employees for all Google services, including your Google Cloud Platform (GCP) organization. What should you do?

A. Use Google Cloud Directory Sync (GCDS) to synchronize users into Cloud Identity.
B. Use the cloud Identity APIs and write a script to synchronize users to Cloud Identity.
C. Export users from Active Directory as a CSV and import them to Cloud Identity via the Admin Console.
D. Ask each employee to create a Google account using self signu
E. Require that each employee use their company email address and password.

Answer: A

NEW QUESTION 133
You have an application running in Google Kubernetes Engine (GKE) with cluster autoscaling enabled. The application exposes a TCP endpoint. There are several replicas of this application. You have a Compute Engine instance in the same region, but in another Virtual Private Cloud (VPC), called gce-network, that has no overlapping IP ranges with the first VPC. This instance needs to connect to the application on GKE. You want to minimize effort. What should you do?

A. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.2. Set the service's externalTrafficPolicy to Cluster.3. Configure the Compute Engine instance to use the address of the load balancer that has been created.
B. 1. In GKE, create a Service of type NodePort that uses the application's Pods as backend.2. Create a Compute Engine instance called proxy with 2 network interfaces, one in each VPC.3. Use iptables on this instance to forward traffic from gce-network to the GKE nodes.4. Configure the Compute Engine instance to use the address of proxy in gce-network as endpoint.
C. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.2. Add an annotation to this service: cloud.google.com/load-balancer-type: Internal3. Peer the two VPCs together.4. Configure the Compute Engine instance to use the address of the load balancer that has been created.
D. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.2. Add a Cloud Armor Security Policy to the load balancer that whitelists the internal IPs of the MIG's instances.3. Configure the Compute Engine instance to use the address of the load balancer that has been created.

Answer: A

NEW QUESTION 136
You have downloaded and installed the gcloud command line interface (CLI) and have authenticated with your Google Account. Most of your Compute Engine instances in your project run in the europe-west1-d zone. You want to avoid having to specify this zone with each CLI command when managing these instances. What should you do?

A. Set the europe-west1-d zone as the default zone using the gcloud config subcommand.
B. In the Settings page for Compute Engine under Default location, set the zone to europe–west1-d.
C. In the CLI installation directory, create a file called default.conf containing zone=europe–west1–d.
D. Create a Metadata entry on the Compute Engine page with key compute/zone and value europe–west1–d.

Answer: C

NEW QUESTION 137
You have a Dockerfile that you need to deploy on Kubernetes Engine. What should you do?

A. Use kubectl app deploy <dockerfilename>.
B. Use gcloud app deploy <dockerfilename>.
C. Create a docker image from the Dockerfile and upload it to Container Registr
D. Create a Deployment YAML file to point to that imag
E. Use kubectl to create the deployment with that file.
F. Create a docker image from the Dockerfile and upload it to Cloud Storag
G. Create a Deployment YAML file to point to that imag
H. Use kubectl to create the deployment with that file.

Answer: C

NEW QUESTION 140
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 142
Several employees at your company have been creating projects with Cloud Platform and paying for it with their personal credit cards, which the company reimburses. The company wants to centralize all these projects under a single, new billing account. What should you do?

A. Contact cloud-billing@google.com with your bank account details and request a corporate billing account for your company.
B. Create a ticket with Google Support and wait for their call to share your credit card details over the phone.
C. In the Google Platform Console, go to the Resource Manage and move all projects to the root Organization.
D. In the Google Cloud Platform Console, create a new billing account and set up a payment method.

Answer: D

NEW QUESTION 146
Your organization has strict requirements to control access to Google Cloud projects. You need to enable your Site Reliability Engineers (SREs) to approve requests from the Google Cloud support team when an SRE opens a support case. You want to follow Google-recommended practices. What should you do?

A. Add your SREs to roles/iam.roleAdmin role.
B. Add your SREs to roles/accessapproval approver role.
C. Add your SREs to a group and then add this group to roles/iam roleAdmin role.
D. Add your SREs to a group and then add this group to roles/accessapproval approver role.

Answer: A

NEW QUESTION 149
You have created a code snippet that should be triggered whenever a new file is uploaded to a Cloud Storage bucket. You want to deploy this code snippet. What should you do?

A. Use App Engine and configure Cloud Scheduler to trigger the application using Pub/Sub.
B. Use Cloud Functions and configure the bucket as a trigger resource.
C. Use Google Kubernetes Engine and configure a CronJob to trigger the application using Pub/Sub.
D. Use Dataflow as a batch job, and configure the bucket as a data source.

Answer: A

NEW QUESTION 150
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 151
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 154
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 158
You need to create a custom IAM role for use with a GCP service. All permissions in the role must be suitable for production use. You also want to clearly share with your organization the status of the custom role. This will be the first version of the custom role. What should you do?

A. Use permissions in your role that use the ‘supported’ support level for role permission
B. Set the role stage to ALPHA while testing the role permissions.
C. Use permissions in your role that use the ‘supported’ support level for role permission
D. Set the role stage to BETA while testing the role permissions.
E. Use permissions in your role that use the ‘testing’ support level for role permission
F. Set the role stage to ALPHA while testing the role permissions.
G. Use permissions in your role that use the ‘testing’ support level for role permission
H. Set the role stage to BETA while testing the role permissions.

Answer: C

NEW QUESTION 159
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 160
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 163
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 166
You want to find out when users were added to Cloud Spanner Identity Access Management (IAM) roles on your Google Cloud Platform (GCP) project. What should you do in the GCP Console?

A. Open the Cloud Spanner console to review configurations.
B. Open the IAM & admin console to review IAM policies for Cloud Spanner roles.
C. Go to the Stackdriver Monitoring console and review information for Cloud Spanner.
D. Go to the Stackdriver Logging console, review admin activity logs, and filter them for Cloud Spanner IAM roles.

Answer: B

NEW QUESTION 167
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 171
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

NEW QUESTION 174
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 178
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 181
You need to monitor resources that are distributed over different projects in Google Cloud Platform. You want to consolidate reporting under the same Stackdriver Monitoring dashboard. What should you do?

A. Use Shared VPC to connect all projects, and link Stackdriver to one of the projects.
B. For each project, create a Stackdriver accoun
C. In each project, create a service account for that project and grant it the role of Stackdriver Account Editor in all other projects.
D. Configure a single Stackdriver account, and link all projects to the same account.
E. Configure a single Stackdriver account for one of the project
F. In Stackdriver, create a Group and add the other project names as criteria for that Group.

Answer: D

NEW QUESTION 183
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 187
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 188
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 191
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 194
Your company has a single sign-on (SSO) identity provider that supports Security Assertion Markup Language (SAML) integration with service providers. Your company has users in Cloud Identity. You would like users to authenticate using your company’s SSO provider. What should you do?

A. In Cloud Identity, set up SSO with Google as an identity provider to access custom SAML apps.
B. In Cloud Identity, set up SSO with a third-party identity provider with Google as a service provider.
C. Obtain OAuth 2.0 credentials, configure the user consent screen, and set up OAuth 2.0 for Mobile & Desktop Apps.
D. Obtain OAuth 2.0 credentials, configure the user consent screen, and set up OAuth 2.0 for Web Server Applications.

Answer: C

NEW QUESTION 196
.....

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