Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

> All examinations will be up to date.

* 24/7 Quality Support

> We will provide service round the clock.

* 100% Pass Rate

> Our guarantee that you will pass the exam.

* Unique Gurantee

> If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1
You need to grant access for three users so that they can view and edit table data on a Cloud Spanner instance. What should you do?

A. Run gcloud iam roles describe roles/spanner.databaseUse
B. Add the users to the role.
C. Run gcloud iam roles describe roles/spanner.databaseUse
D. Add the users to a new grou
E. Add the group to the role.
F. Run gcloud iam roles describe roles/spanner.viewer --project my-projec
G. Add the users to the role.
H. Run gcloud iam roles describe roles/spanner.viewer --project my-projec
I. Add the users to a new group.Add the group to the role.

Answer: A

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 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 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 are working with a user to set up an application in a new VPC behind a firewall. The user is concerned about data egress. You want to configure the fewest open egress ports. What should you do?

A. Set up a low-priority (65534) rule that blocks all egress and a high-priority rule (1000) that allows only the appropriate ports.
B. Set up a high-priority (1000) rule that pairs both ingress and egress ports.
C. Set up a high-priority (1000) rule that blocks all egress and a low-priority (65534) rule that allows only the appropriate ports.
D. Set up a high-priority (1000) rule to allow the appropriate ports.

Answer: C

NEW QUESTION 6
You need to manage a third-party application that will run on a Compute Engine instance. Other Compute Engine instances are already running with default configuration. Application installation files are hosted on Cloud Storage. You need to access these files from the new instance without allowing other virtual machines (VMs) to access these files. What should you do?

A. Create the instance with the default Compute Engine service account Grant the service account permissions on Cloud Storage.
B. Create the instance with the default Compute Engine service account Add metadata to the objects on Cloud Storage that matches the metadata on the new instance.
C. Create a new service account and assign this service account to the new instance Grant the service account permissions on Cloud Storage.
D. Create a new service account and assign this service account to the new instance Add metadata to the objects on Cloud Storage that matches the metadata on the new instance.

Answer: C

NEW QUESTION 7

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 threshol
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 threshol
G. Create a Cloud Function that listens to HTTP and resizes Spanner resources accordingly.

Answer: D

NEW QUESTION 8
Your company has an existing GCP organization with hundreds of projects and a billing account. Your company recently acquired another company that also has hundreds of projects and its own billing account. You would like to consolidate all GCP costs of both GCP organizations onto a single invoice. You would like to consolidate all costs as of tomorrow. What should you do?

A. Link the acquired company's projects to your company's billing account.
B. Configure the acquired company's billing account and your company's billing account to export the billing data into the same BigQuery dataset.
C. Migrate the acquired company's projects into your company's GCP organizatio
D. Link the migrated projects to your company's billing account.
E. Create a new GCP organization and a new billing accoun
F. Migrate the acquired company's projects and your company's projects into the new GCP organization and link the projects to the new billing account.

Answer: D

NEW QUESTION 9
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 10
You are using Data Studio to visualize a table from your data warehouse that is built on top of BigQuery. Data is appended to the data warehouse during the day. At night, the daily summary is recalculated by overwriting the table. You just noticed that the charts in Data Studio are broken, and you want to analyze the problem. What should you do?

A. Use the BigQuery interface to review the nightly Job and look for any errors
B. Review the Error Reporting page in the Cloud Console to find any errors.
C. In Cloud Logging create a filter for your Data Studio report
D. Use Cloud Debugger to find out why the data was not refreshed correctly

Answer: D

NEW QUESTION 10
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 14
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 15
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 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 20
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 21
Your company uses BigQuery for data warehousing. Over time, many different business units in your company have created 1000+ datasets across hundreds of projects. Your CIO wants you to examine all
datasets to find tables that contain an employee_ssn column. You want to minimize effort in performing this task. What should you do?

A. Go to Data Catalog and search for employee_ssn in the search box.
B. Write a shell script that uses the bq command line tool to loop through all the projects in your organization.
C. Write a script that loops through all the projects in your organization and runs a query on INFORMATION_SCHEMA.COLUMNS view to find the employee_ssn column.
D. Write a Cloud Dataflow job that loops through all the projects in your organization and runs a query on INFORMATION_SCHEMA.COLUMNS view to find employee_ssn column.

Answer: D

NEW QUESTION 23
You are deploying an application to App Engine. You want the number of instances to scale based on request rate. You need at least 3 unoccupied instances at all times. Which scaling type should you use?

A. Manual Scaling with 3 instances.
B. Basic Scaling with min_instances set to 3.
C. Basic Scaling with max_instances set to 3.
D. Automatic Scaling with min_idle_instances set to 3.

Answer: D

NEW QUESTION 28
You built an application on your development laptop that uses Google Cloud services. Your application uses Application Default Credentials for authentication and works fine on your development laptop. You want to migrate this application to a Compute Engine virtual machine (VM) and set up authentication using Google-recommended practices and minimal changes. What should you do?

A. Assign appropriate access for Google services to the service account used by the Compute Engine VM.

B. Create a service account with appropriate access for Google services, and configure the application to use this account.
C. Store credentials for service accounts with appropriate access for Google services in a config file, and deploy this config file with your application.
D. Store credentials for your user account with appropriate access for Google services in a config file, and deploy this config file with your application.

Answer: B

NEW QUESTION 31

Your development team needs a new Jenkins server for their project. You need to deploy the server using the fewest steps possible. What should you do?

A. Download and deploy the Jenkins Java WAR to App Engine Standard.
B. Create a new Compute Engine instance and install Jenkins through the command line interface.
C. Create a Kubernetes cluster on Compute Engine and create a deployment with the Jenkins Docker image.
D. Use GCP Marketplace to launch the Jenkins solution.

Answer: D

NEW QUESTION 34

You have an object in a Cloud Storage bucket that you want to share with an external company. The object contains sensitive data. You want access to the content to be removed after four hours. The external company does not have a Google account to which you can grant specific user-based access privileges. You want to use the most secure method that requires the fewest steps. What should you do?

A. Create a signed URL with a four-hour expiration and share the URL with the company.
B. Set object access to 'public' and use object lifecycle management to remove the object after four hours.
C. Configure the storage bucket as a static website and furnish the object's URL to the compan
D. Delete the object from the storage bucket after four hours.
E. Create a new Cloud Storage bucket specifically for the external company to acces
F. Copy the object to that bucke
G. Delete the bucket after four hours have passed.

Answer: A

NEW QUESTION 38

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

Your company runs its Linux workloads on Compute Engine instances. Your company will be working with a new operations partner that does not use Google Accounts. You need to grant access to the instances to your operations partner so they can maintain the installed tooling. What should you do?

A. Enable Cloud IAP for the Compute Engine instances, and add the operations partner as a Cloud IAP Tunnel User.
B. Tag all the instances with the same network ta
C. Create a firewall rule in the VPC to grant TCP access on port 22 for traffic from the operations partner to instances with the network tag.
D. Set up Cloud VPN between your Google Cloud VPC and the internal network of the operations partner.
E. Ask the operations partner to generate SSH key pairs, and add the public keys to the VM instances.

Answer: B

NEW QUESTION 42

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 43

You are using Deployment Manager to create a Google Kubernetes Engine cluster. Using the same Deployment Manager deployment, you also want to create a DaemonSet in the kube-system namespace of the cluster. You want a solution that uses the fewest possible services. What should you do?

A. Add the cluster's API as a new Type Provider in Deployment Manager, and use the new type to create the DaemonSet.
B. Use the Deployment Manager Runtime Configurator to create a new Config resource that contains the DaemonSet definition.
C. With Deployment Manager, create a Compute Engine instance with a startup script that uses kubectl to create the DaemonSet.
D. In the cluster's definition in Deployment Manager, add a metadata that has kube-system as key and the DaemonSet manifest as value.

Answer: C

NEW QUESTION 47
You created a Google Cloud Platform project with an App Engine application inside the project. You initially configured the application to be served from the us-central region. Now you want the application to be served from the asia-northeast1 region. What should you do?

A. Change the default region property setting in the existing GCP project to asia-northeast1.
B. Change the region property setting in the existing App Engine application from us-central to asia-northeast1.
C. Create a second App Engine application in the existing GCP project and specify asia-northeast1 as the region to serve your application.
D. Create a new GCP project and create an App Engine application inside this new projec
E. Specify asia-northeast1 as the region to serve your application.

Answer: C

NEW QUESTION 52
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 54
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 57
You are developing a new web application that will be deployed on Google Cloud Platform. As part of your release cycle, you want to test updates to your application on a small portion of real user traffic. The majority of the users should still be directed towards a stable version of your application. What should you do?

A. Deploy me application on App Engine For each update, create a new version of the same service Configure traffic splitting to send a small percentage of traffic to the new version
B. Deploy the application on App Engine For each update, create a new service Configure traffic splitting to send a small percentage of traffic to the new service.
C. Deploy the application on Kubernetes Engine For a new release, update the deployment to use the new version
D. Deploy the application on Kubernetes Engine For a now release, create a new deployment for the new version Update the service e to use the now deployment.

Answer: A

NEW QUESTION 60
Every employee of your company has a Google account. Your operational team needs to manage a large number of instances on Compute Engine. Each member of this team needs only administrative access to the servers. Your security team wants to ensure that the deployment of credentials is operationally efficient and must be able to determine who accessed a given instance. What should you do?

A. Generate a new SSH key pai
B. Give the private key to each member of your tea
C. Configure the public key in the metadata of each instance.
D. Ask each member of the team to generate a new SSH key pair and to send you their public ke
E. Use a configuration management tool to deploy those keys on each instance.
F. Ask each member of the team to generate a new SSH key pair and to add the public key to their Google accoun
G. Grant the "compute.osAdminLogin" role to the Google group corresponding to this team.
H. Generate a new SSH key pai
I. Give the private key to each member of your tea
J. Configure the public key as a project-wide public SSH key in your Cloud Platform project and allow project-wide public SSH keys on each instance.

Answer: C

Explanation:
https://cloud.google.com/compute/docs/instances/managing-instance-access

NEW QUESTION 65
You've deployed a microservice called myapp1 to a Google Kubernetes Engine cluster using the YAML file specified below:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: myapp1-deployment
spec:
  selector:
    matchLabels:
      app: myapp1
  replicas: 2
  template:
    metadata:
      labels:
        app: myapp1
    spec:
      containers:
      - name: main-container
        image: gcr.io/my-company-repo/myapp1:1.4
        env:
        - name: DB_PASSWORD
          value: "t0ugh2guess!"
        ports:
        - containerPort: 8080
```

You need to refactor this configuration so that the database password is not stored in plain text. You want to follow Google-recommended practices. What should you do?

A. Store the database password inside the Docker image of the container, not in the YAML file.
B. Store the database password inside a Secret objec
C. Modify the YAML file to populate the DB_PASSWORD environment variable from the Secret.
D. Store the database password inside a ConfigMap objec
E. Modify the YAML file to populate the DB_PASSWORD environment variable from the ConfigMap.
F. Store the database password in a file inside a Kubernetes persistent volume, and use a persistent volume claim to mount the volume to the container.

Answer: C

NEW QUESTION 70
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 73
You have a Google Cloud Platform account with access to both production and development projects. You need to create an automated process to list all compute instances in development and production projects on a daily basis. What should you do?

A. Create two configurations using gcloud confi
B. Write a script that sets configurations as active, individuall
C. For each configuration, use gcloud compute instances list to get a list of compute resources.
D. Create two configurations using gsutil confi
E. Write a script that sets configurations as active,individuall
F. For each configuration, use gsutil compute instances list to get a list of compute resources.
G. Go to Cloud Shell and export this information to Cloud Storage on a daily basis.
H. Go to GCP Console and export this information to Cloud SQL on a daily basis.

Answer: A

NEW QUESTION 78
You have a large 5-TB AVRO file stored in a Cloud Storage bucket. Your analysts are proficient only in SQL and need access to the data stored in this file. You want to find a cost-effective way to complete their request as soon as possible. What should you do?

A. Load data in Cloud Datastore and run a SQL query against it.
B. Create a BigQuery table and load data in BigQuer
C. Run a SQL query on this table and drop this table after you complete your request.
D. Create external tables in BigQuery that point to Cloud Storage buckets and run a SQL query on these external tables to complete your request.
E. Create a Hadoop cluster and copy the AVRO file to NDFS by compressing i
F. Load the file in a hive table and provide access to your analysts so that they can run SQL queries.

Answer: C

NEW QUESTION 82
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 86
You need to update a deployment in Deployment Manager without any resource downtime in the deployment. Which command should you use?

A. gcloud deployment-manager deployments create --config <deployment-config-path>
B. gcloud deployment-manager deployments update --config <deployment-config-path>
C. gcloud deployment-manager resources create --config <deployment-config-path>
D. gcloud deployment-manager resources update --config <deployment-config-path>

Answer: B

NEW QUESTION 89
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 92
You are running multiple microservices in a Kubernetes Engine cluster. One microservice is rendering images. The microservice responsible for the image rendering requires a large amount of CPU time compared to the memory it requires. The other microservices are workloads that are optimized for n1-standard machine types. You need to optimize your cluster so that all workloads are using resources as efficiently as possible. What should you do?

A. Assign the pods of the image rendering microservice a higher pod priority than the older microservices
B. Create a node pool with compute-optimized machine type nodes for the image rendering microservice Use the node pool with general-purposemachine type nodes for the other microservices
C. Use the node pool with general-purpose machine type nodes for lite mage rendering microservice Create a nodepool with compute-optimized machine type nodes for the other microservices
D. Configure the required amount of CPU and memory in the resource requests specification of the image rendering microservice deployment Keep the resource requests for the other microservices at the default

Answer: B

NEW QUESTION 97
You have developed a containerized web application that will serve Internal colleagues during business hours. You want to ensure that no costs are incurred outside of the hours the application is used. You have just created a new Google Cloud project and want to deploy the application. What should you do?

A. Deploy the container on Cloud Run for Anthos, and set the minimum number of instances to zero
B. Deploy the container on Cloud Run (fully managed), and set the minimum number of instances to zero.
C. Deploy the container on App Engine flexible environment with autoscalin
D. and set the value min_instances to zero in the app yaml
E. Deploy the container on App Engine flexible environment with manual scaling, and set the value instances to zero in the app yaml

Answer: C

NEW QUESTION 101
You are assisting a new Google Cloud user who just installed the Google Cloud SDK on their VM. The server needs access to Cloud Storage. The user wants your help to create a new storage bucket. You need to make this change in multiple environments. What should you do?

A. Use a Deployment Manager script to automate creating storage buckets in an appropriate region
B. Use a local SSD to improve performance of the VM for the targeted workload
C. Use the gsutii command to create a storage bucket in the same region as the VM
D. Use a Persistent Disk SSD in the same zone as the VM to improve performance of the VM

Answer: A

NEW QUESTION 104
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 are the team lead of a group of 10 developers. You provided each developer with an individual Google Cloud Project that they can use as their personal sandbox to experiment with different Google Cloud solutions. You want to be notified if any of the developers are spending above $500 per month on their sandbox environment. What should you do?

A. Create a single budget for all projects and configure budget alerts on this budget.
B. Create a separate billing account per sandbox project and enable BigQuery billing export
C. Create a Data Studio dashboard to plot the spending per billing account.
D. Create a budget per project and configure budget alerts on all of these budgets.
E. Create a single billing account for all sandbox projects and enable BigQuery billing export
F. Create a Data Studio dashboard to plot the spending per project.

Answer: C

NEW QUESTION 111
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 112
You have an application that looks for its licensing server on the IP 10.0.3.21. You need to deploy the licensing server on Compute Engine. You do not want to change the configuration of the application and want the application to be able to reach the licensing server. What should you do?

A. Reserve the IP 10.0.3.21 as a static internal IP address using gcloud and assign it to the licensing server.
B. Reserve the IP 10.0.3.21 as a static public IP address using gcloud and assign it to the licensing server.
C. Use the IP 10.0.3.21 as a custom ephemeral IP address and assign it to the licensing server.
D. Start the licensing server with an automatic ephemeral IP address, and then promote it to a static internal IP address.

Answer: A

NEW QUESTION 116
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 118
You have a web application deployed as a managed instance group. You have a new version of the application to gradually deploy. Your web application is currently receiving live web traffic. You want to ensure that the available capacity does not decrease during the deployment. What should you do?

A. Perform a rolling-action start-update with maxSurge set to 0 and maxUnavailable set to 1.
B. Perform a rolling-action start-update with maxSurge set to 1 and maxUnavailable set to 0.
C. Create a new managed instance group with an updated instance templat
D. Add the group to the backend service for the load balance
E. When all instances in the new managed instance group are healthy, delete the old managed instance group.
F. Create a new instance template with the new application versio
G. Update the existing managed instance group with the new instance templat
H. Delete the instances in the managed instance group to allow the managed instance group to recreate the instance using the new instance template.

Answer: B

NEW QUESTION 119
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 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 123
You have an application that uses Cloud Spanner as a database backend to keep current state information about users. Cloud Bigtable logs all events triggered by users. You export Cloud Spanner data to Cloud Storage during daily backups. One of your analysts asks you to join data from Cloud Spanner and Cloud Bigtable for specific users. You want to complete this ad hoc request as efficiently as possible. What should you do?

A. Create a dataflow job that copies data from Cloud Bigtable and Cloud Storage for specific users.
B. Create a dataflow job that copies data from Cloud Bigtable and Cloud Spanner for specific users.
C. Create a Cloud Dataproc cluster that runs a Spark job to extract data from Cloud Bigtable and Cloud Storage for specific users.
D. Create two separate BigQuery external tables on Cloud Storage and Cloud Bigtabl
E. Use the BigQuery console to join these tables through user fields, and apply appropriate filters.

Answer: D

NEW QUESTION 127
You have production and test workloads that you want to deploy on Compute Engine. Production VMs need to be in a different subnet than the test VMs. All the VMs must be able to reach each other over internal IP without creating additional routes. You need to set up VPC and the 2 subnets. Which configuration meets these requirements?

A. Create a single custom VPC with 2 subnet
B. Create each subnet in a different region and with a different CIDR range.
C. Create a single custom VPC with 2 subnet
D. Create each subnet in the same region and with the same CIDR range.
E. Create 2 custom VPCs, each with a single subne
F. Create each subnet is a different region and with a different CIDR range.
G. Create 2 custom VPCs, each with a single subne
H. Create each subnet in the same region and with the same CIDR range.

Answer: A

NEW QUESTION 129
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 134
&. You are working for a hospital that stores Its medical images in an on-premises data room. The hospital wants to use Cloud Storage for archival storage of these images. The hospital wants an automated process to upload any new medical images to Cloud Storage. You need to design and implement a solution. What should you do?

A. Deploy a Dataflow job from the batch template "Datastore lo Cloud Storage" Schedule the batch job on the desired interval
B. In the Cloud Console, go to Cloud Storage Upload the relevant images to the appropriate bucket
C. Create a script that uses the gsutil command line interface to synchronize the on-premises storage with Cloud Storage Schedule the script as a cron job
D. Create a Pub/Sub topic, and enable a Cloud Storage trigger for the Pub/Sub topi
E. Create an application that sends all medical images to the Pub/Sub lope

Answer: A

NEW QUESTION 138
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 141

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 145

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 150

You are operating a Google Kubernetes Engine (GKE) cluster for your company where different teams can run non-production workloads. Your Machine Learning (ML) team needs access to Nvidia Tesla P100 GPUs to train their models. You want to minimize effort and cost. What should you do?

A. Ask your ML team to add the "accelerator: gpu" annotation to their pod specification.
B. Recreate all the nodes of the GKE cluster to enable GPUs on all of them.
C. Create your own Kubernetes cluster on top of Compute Engine with nodes that have GPU
D. Dedicate this cluster to your ML team.
E. Add a new, GPU-enabled, node pool to the GKE cluste
F. Ask your ML team to add the cloud.google.com/gke -accelerator: nvidia-tesla-p100 nodeSelector to their pod specification.

Answer: B

NEW QUESTION 152

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 155

You want to send and consume Cloud Pub/Sub messages from your App Engine application. The Cloud Pub/Sub API is currently disabled. You will use a service account to authenticate your application to the API. You want to make sure your application can use Cloud Pub/Sub. What should you do?

A. Enable the Cloud Pub/Sub API in the API Library on the GCP Console.
B. Rely on the automatic enablement of the Cloud Pub/Sub API when the Service Account accesses it.
C. Use Deployment Manager to deploy your applicatio
D. Rely on the automatic enablement of all APIs used by the application being deployed.
E. Grant the App Engine Default service account the role of Cloud Pub/Sub Admi
F. Have your application enable the API on the first connection to Cloud Pub/Sub.

Answer: A

NEW QUESTION 159
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 160
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 161
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 166
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 168
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 170
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 172
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 176
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 181
You need to enable traffic between multiple groups of Compute Engine instances that are currently running two different GCP projects. Each group of Compute Engine instances is running in its own VPC. What should you do?

A. Verify that both projects are in a GCP Organizatio
B. Create a new VPC and add all instances.
C. Verify that both projects are in a GCP Organizatio
D. Share the VPC from one project and request that the Compute Engine instances in the other project use this shared VPC.
E. Verify that you are the Project Administrator of both project
F. Create two new VPCs and add all instances.
G. Verify that you are the Project Administrator of both project
H. Create a new VPC and add all instances.

Answer: B

NEW QUESTION 184
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 189
You have been asked to create robust Virtual Private Network (VPN) connectivity between a new Virtual Private Cloud (VPC) and a remote site. Key requirements include dynamic routing, a shared address space of 10.19.0.1/22, and no overprovisioning of tunnels during a failover event. You want to follow Google-recommended practices to set up a high availability Cloud VPN. What should you do?

A. Use a custom mode VPC network, configure static routes, and use active/passive routing
B. Use an automatic mode VPC network, configure static routes, and use active/active routing
C. Use a custom mode VPC network use Cloud Router border gateway protocol (86P) routes, and use active/passive routing
D. Use an automatic mode VPC network, use Cloud Router border gateway protocol (BGP) routes and configure policy-based routing

Answer: D

NEW QUESTION 193
......

Relate Links

100% Pass Your Associate-Cloud-Engineer Exam with Exambible Prep Materials

https://www.exambible.com/Associate-Cloud-Engineer-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/