



Vendor: Google

Exam Code: Associate-Cloud-Engineer

Exam Name: Associate Cloud Engineer

Version: DEMO

QUESTION 1

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

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

Answer: A

Explanation:

Google Cloud Deployment Manager allows you to specify all the resources needed for your application in a declarative format using yaml. You can also use Python or Jinja2 templates to parameterize the configuration and allow reuse of common deployment paradigms such as a load balanced, auto-scaled instance group. Treat your configuration as code and perform repeatable deployments.

<https://cloud.google.com/deployment-manager>

QUESTION 2

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 read. Then convert that bytes estimate to dollars using the Pricing Calculator.
- C. Use the command line to run a dry run query to estimate the number of bytes returned. Then convert that bytes estimate to dollars using the Pricing Calculator.
- D. Run a select count (*) to get an idea of how many records your query will look through. Then convert that number of rows to dollars using the Pricing Calculator.

Answer: B

Explanation:

On-demand pricing

Under on-demand pricing, BigQuery charges for queries by using one metric: the number of bytes processed (also referred to as bytes read). You are charged for the number of bytes processed whether the data is stored in BigQuery or in an external data source such as Cloud Storage, Drive, or Cloud Bigtable. On-demand pricing is based solely on usage.

https://cloud.google.com/bigquery/pricing#on_demand_pricing

QUESTION 3

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: A

Explanation:

As per google best practices it is recommended to use predefined roles and create groups to control access to multiple users with same responsibility.

<https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#groups-and-service-accounts> .

QUESTION 4

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 project. Specify asia-northeast1 as the region to serve your application.

Answer: D

Explanation:

Each Cloud project can contain only a single App Engine application, and once created you cannot change the location of your App Engine application.

<https://cloud.google.com/appengine/docs/flexible/nodejs/managing-projects-apps-billing#create>

QUESTION 5

You create a new Google Kubernetes Engine (GKE) cluster and want to make sure that it always runs a supported and stable version of Kubernetes. What should you do?

- A. Enable the Node Auto-Repair feature for your GKE cluster.
- B. Enable the Node Auto-Upgrades feature for your GKE cluster.
- C. Select the latest available cluster version for your GKE cluster.
- D. Select "Container-Optimized OS (cos)" as a node image for your GKE cluster.

Answer: B

Explanation:

Creating or upgrading a cluster by specifying the version as <latest> does not provide automatic upgrades. Enable automatic node upgrades to ensure that the nodes in your cluster up to date with the latest stable version.

<https://cloud.google.com/kubernetes-engine/versioning-and-upgrades>

QUESTION 6

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

Explanation:

https://cloud.google.com/storage/docs/gsutil/commands/cp#parallel-composite-uploads_1

Warning: Parallel composite uploads should not be used with NEARLINE, COLDLINE, or ARCHIVE storage class buckets, because doing so incurs an early deletion charge for each component object.

Warning: Parallel composite uploads should not be used in buckets that have a retention policy, because the component pieces cannot be deleted until each has met the bucket's minimum retention period.

QUESTION 7

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

Answer: C

Explanation:

For example, batch processing jobs can run on preemptible instances. If some of those instances stop during processing, the job slows but does not completely stop. Preemptible instances complete your batch processing tasks without placing additional workload on your existing instances and without requiring you to pay full price for additional normal instances.

<https://cloud.google.com/compute/docs/instances/preemptible>

QUESTION 8

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 subnets.
Create each subnet in a different region and with a different CIDR range.
- B. Create a single custom VPC with 2 subnets.
Create each subnet in the same region and with the same CIDR range.
- C. Create 2 custom VPCs, each with a single subnet.
Create each subnet in a different region and with a different CIDR range.
- D. Create 2 custom VPCs, each with a single subnet.
Create each subnet in the same region and with the same CIDR range.

Answer: A

Explanation:

Primary and secondary ranges for subnets cannot overlap with any allocated range, any primary or secondary range of another subnet in the same network, or any IP ranges of subnets in peered networks.

<https://cloud.google.com/vpc/docs/using-vpc#subnet-rules>

QUESTION 9

Your company has a Google Cloud Platform project that uses BigQuery for data warehousing. Your data science team changes frequently and has few members. You need to allow members of this team to perform queries.

You want to follow Google-recommended practices. What should you do?

- A. 1. Create an IAM entry for each data scientist's user account.
2. Assign the BigQuery jobUser role to the group.
- B. 1. Create an IAM entry for each data scientist's user account.
2. Assign the BigQuery dataViewer user role to the group.
- C. 1. Create a dedicated Google group in Cloud Identity.
2. Add each data scientist's user account to the group.
3. Assign the BigQuery jobUser role to the group.
- D. 1. Create a dedicated Google group in Cloud Identity.
2. Add each data scientist's user account to the group.
3. Assign the BigQuery dataViewer user role to the group.

Answer: C

Explanation:

Provides permissions to run jobs, including queries, within the project. This role can check the existence of all jobs, enumerate their own jobs, and cancel their own jobs.

<https://cloud.google.com/bigquery/docs/access-control>

QUESTION 10

You want to configure autohealing for network load balancing for a group of Compute Engine instances that run in multiple zones, using the fewest possible steps. You need to configure re-creation of VMs if they are unresponsive after 3 attempts of 10 seconds each. What should you do?

- A. Create an HTTP load balancer with a backend configuration that references an existing instance group.
Set the health check to healthy (HTTP)
- B. Create an HTTP load balancer with a backend configuration that references an existing instance group.
Define a balancing mode and set the maximum RPS to 10.
- C. Create a managed instance group. Set the Autohealing health check to healthy (HTTP)
- D. Create a managed instance group. Verify that the autoscaling setting is on.

Answer: C

Explanation:

Use separate health checks for load balancing and for autohealing. Health checks for load balancing detect unresponsive instances and direct traffic away from them. Health checks for autohealing detect and recreate failed instances, so they should be less aggressive than load balancing health checks. Using the same health check for these services would remove the distinction between unresponsive instances and failed instances, causing unnecessary latency and unavailability for your users.

<https://cloud.google.com/compute/docs/tutorials/high-availability-autohealing>

QUESTION 11

You are using multiple configurations for gcloud. You want to review the configured Kubernetes Engine cluster of an inactive configuration using the fewest possible steps. What should you do?

- A. Use gcloud config configurations describe to review the output.
- B. Use gcloud config configurations activate and gcloud config list to review the output.
- C. Use kubectl config get-contexts to review the output.
- D. Use kubectl config use-context and kubectl config view to review the output.

Answer: D

Explanation:

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

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

QUESTION 12

You have a single binary application that you want to run on Google Cloud Platform. You decided to automatically scale the application based on underlying infrastructure CPU usage. Your organizational policies require you to use virtual machines directly. You need to ensure that the application scaling is operationally efficient and completed as quickly as possible. What should you do?

- A. Create a Google Kubernetes Engine cluster, and use horizontal pod autoscaling to scale the application.
- B. Create an instance template, and use the template in a managed instance group with autoscaling configured.
- C. Create an instance template, and use the template in a managed instance group that scales up and down based on the time of day.
- D. Use a set of third-party tools to build automation around scaling the application up and down, based on Stackdriver CPU usage monitoring.

Answer: B

Explanation:

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

QUESTION 13

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

- A. When creating the VM via the web console, specify the service account under the 'Identity and API Access' section.
- B. Download a JSON Private Key for the service account. On the Project Metadata, add that JSON as the value for the key compute-engine-service-account.
- C. Download a JSON Private Key for the service account. On the Custom Metadata of the VM, add that JSON as the value for the key compute-engine-service-account.
- D. Download a JSON Private Key for the service account.

After creating the VM, ssh into the VM and save the JSON under ~/.gcloud/compute-engine-service-account.json.

Answer: A

Explanation:

Changing the service account and access scopes for an instance

If you want to run the VM as a different identity, or you determine that the instance needs a different set of scopes to call the required APIs, you can change the service account and the access scopes of an existing instance. For example, you can change access scopes to grant access to a new API, or change an instance so that it runs as a service account that you created, instead of the Compute Engine default service account. However, Google recommends that you use the fine-grained IAM policies instead of relying on access scopes to control resource access for the service account.

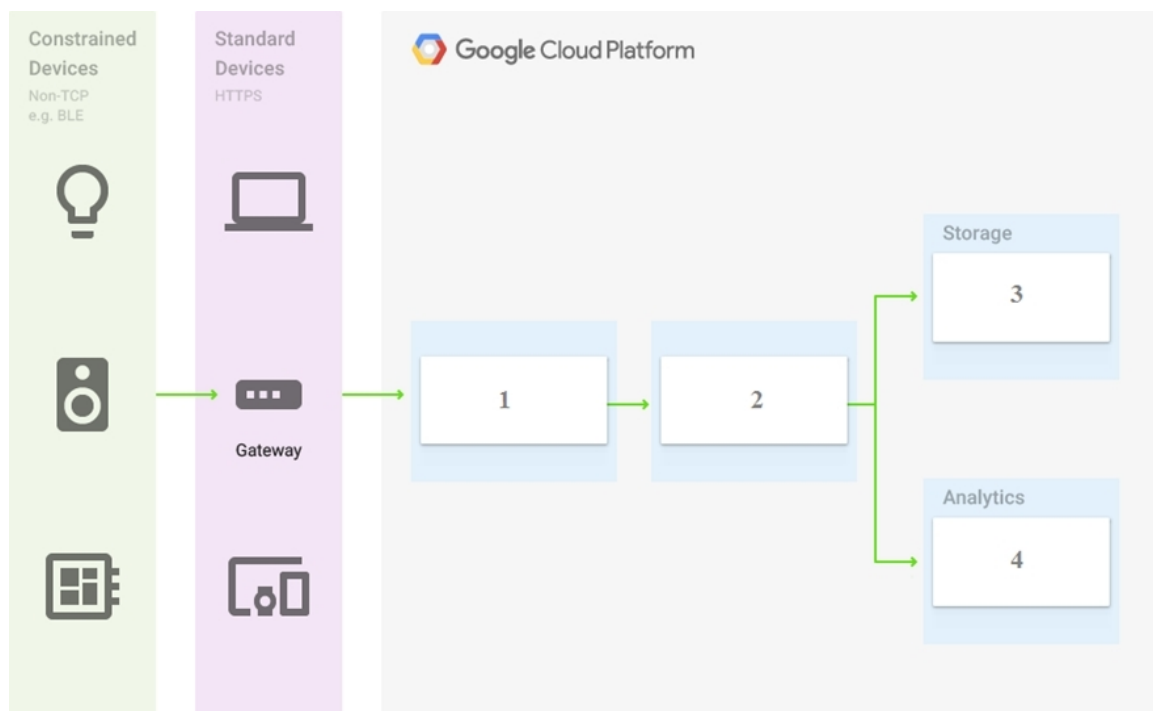
To change an instance's service account and access scopes, the instance must be temporarily stopped. To stop your instance, read the documentation for Stopping an instance. After changing the service account or access scopes, remember to restart the instance. Use one of the following methods to the change service account or access scopes of the stopped instance.

<https://cloud.google.com/compute/docs/access/create-enable-service-accounts-for-instances>

QUESTION 14

You are building a pipeline to process time-series data.

Which Google Cloud Platform services should you put in boxes 1,2,3, and 4?



- A. Cloud Pub/Sub, Cloud Dataflow, Cloud Datastore, BigQuery
- B. Firebase Messages, Cloud Pub/Sub, Cloud Spanner, BigQuery
- C. Cloud Pub/Sub, Cloud Storage, BigQuery, Cloud Bigtable
- D. Cloud Pub/Sub, Cloud Dataflow, Cloud Bigtable, BigQuery

Answer: D

Explanation:

Correct answer is D as Cloud Pub/Sub for data ingestion, Dataflow for data handling and transformation, Bigtable for storage to provide low latency data access and BigQuery for analytics. Whenever we want to process timeseries data look for BigTable.

Also you want to perform analytics in Box 4 ..look for BigQuery

<https://cloud.google.com/blog/products/databases/getting-started-with-time-series-trend-predictions-using-gcp>

QUESTION 15

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 company. Delete the object from the storage bucket after four hours.
- D. Create a new Cloud Storage bucket specifically for the external company to access. Copy the object to that bucket. Delete the bucket after four hours have passed.

Answer: A

Explanation:

Signed URLs are used to give time-limited resource access to anyone in possession of the URL, regardless of whether they have a Google account.

<https://cloud.google.com/storage/docs/access-control/signed-urls>

QUESTION 16

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 Manager 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

Explanation:

Cloud Billing accounts can be used across organizations. Moving a project from one organization to another won't impact billing, and charges will continue against the old billing account. However, organization moves often also include a requirement to move to a new billing account.

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

QUESTION 17

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 pair. Give the private key to each member of your team. Configure the public key in the metadata of each instance.
- B. Ask each member of the team to generate a new SSH key pair and to send you their public key. Use a configuration management tool to deploy those keys on each instance.
- C. Ask each member of the team to generate a new SSH key pair and to add the public key to their Google account. Grant the "compute.osAdminLogin" role to the Google group corresponding to this team.
- D. Generate a new SSH key pair. Give the private key to each member of your team. 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:

We recommend collecting users with the same responsibilities into groups and assigning IAM roles to the groups rather than to individual users. For example, you can create a "data scientist" group and assign appropriate roles to enable interaction with BigQuery and Cloud Storage. When a new data scientist joins your team, you can simply add them to the group and they will inherit the defined permissions. You can create and manage groups through the Admin Console. <https://cloud.google.com/compute/docs/instances/managing-instance-access>

QUESTION 18

You need to create a custom VPC with a single subnet. The subnet's range must be as large as possible.

Which range should you use?

- A. 0.0.0.0/0
- B. 10.0.0.0/8
- C. 172.16.0.0/12
- D. 192.168.0.0/16

Answer: B

Explanation:

The private network range is defined by IETF (Ref: <https://tools.ietf.org/html/rfc1918>) and adhered to by all cloud providers. The supported internal IP Address ranges are:

1. 24-bit block 10.0.0.0/8 (16777216 IP Addresses)
2. 20-bit block 172.16.0.0/12 (1048576 IP Addresses)
3. 16-bit block 192.168.0.0/16 (65536 IP Addresses)

10.0.0.0/8 gives you the most extensive range - 16777216 IP Addresses.

QUESTION 19

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: A

Explanation:

Since both the Billing Account User & Project Billing Manager can link a project to a billing account but clearly states the Project Billing Manager will not have access to resources. Option A is correct, as we don't want the engineering team to link projects to billing account and want only the Finance team. Billing Account User role will help to link projects to the billing account.

<https://cloud.google.com/billing/docs/how-to/billing-accessroles/billing.user> = Link projects to billing accounts

QUESTION 20

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

- A. Create a Cloud Memorystore for Redis instance with 32-GB capacity.
- B. Run it on Compute Engine, and choose a custom instance type with 6 vCPUs and 32 GB of memory.
- C. Package it in a container image, and run it on Kubernetes Engine, using n1-standard-32 instances as nodes.
- D. Run it on Compute Engine, choose the instance type n1-standard-1, and add an SSD persistent disk of 32 GB.

Answer: A

Explanation:

What it's good for Memorystore for Redis provides a fast, in-memory store for use cases that require fast, real-time processing of data. From simple caching use cases to real time analytics, Memorystore for Redis provides the performance you need.

Caching: Cache is an integral part of modern application architectures. Memorystore for Redis provides low latency access and high throughput for heavily accessed data, compared to accessing the data from a disk based backend store. Session management, frequently accessed queries, scripts, and pages are common examples of caching.

https://cloud.google.com/memorystore/docs/redis/redis-overview#what_its_good_for

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.lead2pass.com/all-products.html>



Microsoft



ORACLE



CITRIX



JUNIPER
NETWORKS



EMC²
where information lives

10% Discount Coupon Code: ASTR14