Free Questions for Associate-Cloud-Engineer by certsinside

Shared by Beard on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You just installed the Google Cloud CLI on your new corporate laptop. You need to list the existing instances of your company on Google Cloud. What must you do before you run the gcloud compute instances list command?

Choose 2 answers

Options:

A- Run gcloud auth login, enter your login credentials in the dialog window, and paste the received login token to gcloud CLI.

B- Create a Google Cloud service account, and download the service account key. Place the key file in a folder on your machine where gcloud CLI can find it.

C- Download your Cloud Identity user account key. Place the key file in a folder on your machine where gcloud CLI can find it.

D- Run gcloud config set compute/zone $my_zone to set the default zone for gcloud CLI.

E- Run gcloud config set project $my_project to set the default project for gcloud CLI.

Answer:

A, E

Explanation:

Before you run the gcloud compute instances list command, you need to do two things: authenticate with your user account and set the default project for gcloud CLI.

To authenticate with your user account, you need to run gcloud auth login, enter your login credentials in the dialog window, and paste the received login token to gcloud CLI.This will authorize the gcloud CLI to access Google Cloud resources on your behalf1.

To set the default project for gcloud CLI, you need to run gcloud config set project $my_project, where $my_project is the ID of the project that contains the instances you want to list.This will save you from having to specify the project flag for every gcloud command2.

Option B is not recommended, because using a service account key increases the risk of credential leakage and misuse.It is also not necessary, because you can use your user account to authenticate to the gcloud CLI3. Option C is not correct, because there is no such thing as a Cloud Identity user account key.Cloud Identity is a service that provides identity and access management for Google Cloud users and groups4. Option D is not required, because the gcloud compute instances list command does not depend on the default zone. You can list instances from all zones or filter by a specific zone using the --filter flag.

1: https://cloud.google.com/sdk/docs/authorizing

2: https://cloud.google.com/sdk/gcloud/reference/config/set

3: https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys

4: https://cloud.google.com/identity/docs/overview

: https://cloud.google.com/sdk/gcloud/reference/compute/instances/list

Question 2

Question Type: MultipleChoice

You have deployed an application on a Compute Engine instance. An external consultant needs to access the Linux-based instance. The consultant is connected to your corporate network through a VPN connection, but the consultant has no Google account. What should you do?

Options:

A- Instruct the external consultant to use the gcloud compute ssh command line tool by using Identity-Aware Proxy to access the instance.

B- Instruct the external consultant to use the gcloud compute ssh command line tool by using the public IP address of the instance to access it.

C- Instruct the external consultant to generate an SSH key pair, and request the public key from the consultant.
Add the public key to the instance yourself, and have the consultant access the instance through SSH with their private key.

D- Instruct the external consultant to generate an SSH key pair, and request the private key from the consultant.
Add the private key to the instance yourself, and have the consultant access the instance through SSH with their public key.

Answer:

C

Explanation:

The best option is to instruct the external consultant to generate an SSH key pair, and request the public key from the consultant. Then, add the public key to the instance yourself, and have the consultant access the instance through SSH with their private key. This way, you can grant the consultant access to the instance without requiring a Google account or exposing the instance's public IP address.This option also follows the best practice of using user-managed SSH keys instead of service account keys for SSH access1.

Option A is not feasible because the external consultant does not have a Google account, and therefore cannot use Identity-Aware Proxy (IAP) to access the instance.IAP requires the user to authenticate with a Google account and have the appropriate IAM permissions to access the instance2. Option B is not secure because it exposes the instance's public IP address, which can increase the risk of unauthorized access or attacks. Option D is not correct because it reverses the roles of the public and private keys. The public key should be added to the instance, and the private key should be kept by the consultant.Sharing the private key with anyone else can compromise the security of the SSH connection3.

1: https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys

2: https://cloud.google.com/iap/docs/using-tcp-forwarding

3: https://cloud.google.com/compute/docs/instances/connecting-advanced#sshbetweeninstances

Question 3

Question Type: MultipleChoice

You recently discovered that your developers are using many service account keys during their development process. While you work on a long term improvement, you need to quickly implement a process to enforce short-lived service account credentials in your company. You have the following requirements:

* All service accounts that require a key should be created in a centralized project called pj-sa.

* Service account keys should only be valid for one day.

You need a Google-recommended solution that minimizes cost. What should you do?

Options:

A- Implement a Cloud Run job to rotate all service account keys periodically in pj-sa. Enforce an org policy to deny service account key creation with an exception to pj-sa.

B- Implement a Kubernetes Cronjob to rotate all service account keys periodically. Disable attachment of service accounts to resources in all projects with an exception to pj-sa.

C- Enforce an org policy constraint allowing the lifetime of service account keys to be 24 hours. Enforce an org policy constraint denying service account key creation with an exception on pj-sa.

D- Enforce a DENY org policy constraint over the lifetime of service account keys for 24 hours. Disable attachment of service accounts to resources in all projects with an exception to pj-sa.

Answer:

C

Explanation:

According to the Google Cloud documentation, you can use organization policy constraints to control the creation and expiration of service account keys. The constraints are:

constraints/iam.allowServiceAccountKeyCreation: This constraint allows you to specify which projects or folders can create service account keys. You can set the value totrueorfalse, or use a condition to apply the constraint to specific service accounts. By setting this constraint tofalsefor the organization and adding an exception for the pj-sa project, you can prevent developers from creating service account keys in other projects.

constraints/iam.serviceAccountKeyMaxLifetime: This constraint allows you to specify the maximum lifetime of service account keys. You can set the value to a duration in seconds, such as86400for one day. By setting this constraint to86400for the organization, you can ensure that all service account keys expire after one day.

These constraints are recommended by Google Cloud as best practices to minimize the risk of service account key misuse or compromise. They also help you reduce the cost of managing service account keys, as you do not need to implement a custom solution to rotate or delete them.

1: Associate Cloud Engineer Certification Exam Guide | Learn - Google Cloud

5: Create and delete service account keys - Google Cloud

Organization policy constraints for service accounts

Question 4

Question Type: MultipleChoice

Your company is moving its continuous integration and delivery (CI/CD) pipeline to Compute Engine instances. The pipeline will manage the entire cloud infrastructure through code. How can you ensure that the pipeline has appropriate permissions while your system is following security best practices?

Options:

A- * Add a step for human approval to the CI/CD pipeline before the execution of the infrastructure provisioning.
* Use the human approvals IAM account for the provisioning.

B- * Attach a single service account to the compute instances.
* Add minimal rights to the service account.
* Allow the service account to impersonate a Cloud Identity user with elevated permissions to create, update, or delete resources.

C- * Attach a single service account to the compute instances.
* Add all required Identity and Access Management (IAM) permissions to this service account to create, update, or delete resources

D- * Create multiple service accounts, one for each pipeline with the appropriate minimal Identity and Access Management (IAM) permissions.
* Use a secret manager service to store the key files of the service accounts.
* Allow the CI/CD pipeline to request the appropriate secrets during the execution of the pipeline.

Answer:

B

Explanation:

The best option is to attach a single service account to the compute instances and add minimal rights to the service account. Then, allow the service account to impersonate a Cloud Identity user with elevated permissions to create, update, or delete resources. This way, the service account can use short-lived access tokens to authenticate to Google Cloud APIs without needing to manage service account keys. This option follows the principle of least privilege and reduces the risk of credential leakage and misuse.

Option A is not recommended because it requires human intervention, which can slow down the CI/CD pipeline and introduce human errors. Option C is not secure because it grants all required IAM permissions to a single service account, which can increase the impact of a compromised key. Option D is not cost-effective because it requires creating and managing multiple service accounts and keys, as well as using a secret manager service.

1: https://cloud.google.com/iam/docs/impersonating-service-accounts

2: https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys

3: https://cloud.google.com/iam/docs/understanding-service-accounts

Question 5

Question Type: MultipleChoice

Your continuous integration and delivery (CI/CD) server can't execute Google Cloud actions in a specific project because of permission issues. You need to validate whether the used service account has the appropriate roles in the specific project. What should you do?

Options:

A- Open the Google Cloud console, and run a query to determine which resources this service account can access.

B- Open the Google Cloud console, and run a query of the audit logs to find permission denied errors for this service account.

C- Open the Google Cloud console, and check the organization policies.

D- Open the Google Cloud console, and check the Identity and Access Management (IAM) roles assigned to the service account at the project or inherited from the folder or organization levels.

Answer:

D

Explanation:

This answer is the most effective way to validate whether the service account used by the CI/CD server has the appropriate roles in the specific project. By checking the IAM roles assigned to the service account, you can see which permissions the service account has and

which resources it can access. You can also check if the service account inherits any roles from the folder or organization levels, which may affect its access to the project. You can use the Google Cloud console, the gcloud command-line tool, or the IAM API to view the IAM roles of a service account.

Question 6

Question Type: MultipleChoice

After a recent security incident, your startup company wants better insight into what is happening in the Google Cloud environment. You need to monitor unexpected firewall changes and instance creation. Your company prefers simple solutions. What should you do?

Options:

A- Use Cloud Logging filters to create log-based metrics for firewall and instance actions. Monitor the changes and set up reasonable alerts.

B- Install Kibana on a compute Instance. Create a log sink to forward Cloud Audit Logs filtered for firewalls and compute instances to Pub/Sub. Target the Pub/Sub topic to push messages to the Kibana instance. Analyze the logs on Kibana in real time.

C- Turn on Google Cloud firewall rules logging, and set up alerts for any insert, update, or delete events.

D- Create a log sink to forward Cloud Audit Logs filtered for firewalls and compute instances to Cloud Storage.

Use BigQuery to periodically analyze log events in the storage bucket.

Answer:

A

Explanation:

This answer is the simplest and most effective way to monitor unexpected firewall changes and instance creation in Google Cloud. Cloud Logging filters allow you to specify the criteria for the log entries that you want to view or export. You can use the Logging query language to write filters based on the LogEntry fields, such as resource.type, severity, or protoPayload.methodName. For example, you can filter for firewall-related events by using the following query:

resource.type="gce_subnetwork" logName="projects/PROJECT_ID/logs/compute.googleapis.com%2Ffirewall"

You can filter for instance-related events by using the following query:

resource.type="gce_instance" logName="projects/PROJECT_ID/logs/compute.googleapis.com%2Factivity_log"

You can create log-based metrics from these filters to measure the rate or count of log entries that match the filter. Log-based metrics can be used to create charts and dashboards in Cloud Monitoring, or to set up alerts based on the metric values. For example, you can create an alert policy that triggers when the log-based metric for firewall changes exceeds a certain threshold in a given time interval. This way, you can get notified of any unexpected or malicious changes to your firewall rules.

Option B is incorrect because it is unnecessarily complex and costly. Installing Kibana on a compute instance requires additional configuration and maintenance. Creating a log sink to forward Cloud Audit Logs to Pub/Sub also incurs additional charges for the

Pub/Sub service. Analyzing the logs on Kibana in real time may not be feasible or efficient, as it requires constant monitoring and manual intervention.

Option C is incorrect because Google Cloud firewall rules logging is a different feature from Cloud Audit Logs. Firewall rules logging allows you to audit, verify, and analyze the effects of your firewall rules by creating connection records for each rule that applies to traffic. However, firewall rules logging does not log the insert, update, or delete events for the firewall rules themselves. Those events are logged by Cloud Audit Logs, which record the administrative activities in your Google Cloud project.

Option D is incorrect because it is not a real-time solution. Creating a log sink to forward Cloud Audit Logs to Cloud Storage requires additional storage space and charges. Using BigQuery to periodically analyze log events in the storage bucket also incurs additional costs for the BigQuery service. Moreover, this option does not provide any alerting mechanism to notify you of any unexpected or malicious changes to your firewall rules or instances.

Question 7

Question Type: **MultipleChoice**

You are in charge of provisioning access for all Google Cloud users in your organization. Your company recently acquired a startup company that has their own Google Cloud organization. You need to ensure that your Site Reliability Engineers (SREs) have the same project permissions in the startup company's organization as in your own organization. What should you do?

Options:

A- In the Google Cloud console for your organization, select Create role from selection, and choose destination as the startup company's organization

B- In the Google Cloud console for the startup company, select Create role from selection and choose source as the startup company's Google Cloud organization.

C- Use the gcloud iam roles copy command, and provide the Organization ID of the startup company's Google Cloud Organization as the destination.

D- Use the gcloud iam roles copy command, and provide the project IDs of all projects in the startup company s organization as the destination.

Answer:

C

Explanation:

https://cloud.google.com/architecture/best-practices-vpc-design#shared-service Cloud VPN is another alternative. Because Cloud VPN establishes reachability through managed IPsec tunnels, it doesn't have the aggregate limits of VPC Network Peering. Cloud VPN uses a VPN Gateway for connectivity and doesn't consider the aggregate resource use of the IPsec peer. The drawbacks of Cloud VPN include increased costs (VPN tunnels and traffic egress), management overhead required to maintain tunnels, and the performance overhead of IPsec.

Question 8

Question Type: MultipleChoice

You have an on-premises data analytics set of binaries that processes data files in memory for about 45 minutes every midnight. The sizes of those data files range from 1 gigabyte to 16 gigabytes. You want to migrate this application to Google Cloud with minimal effort and cost. What should you do?

Options:

A- Upload the code to Cloud Functions. Use Cloud Scheduler to start the application.

B- Create a container for the set of binaries. Use Cloud Scheduler to start a Cloud Run job for the container.

C- Create a container for the set of binaries Deploy the container to Google Kubernetes Engine (GKE) and use the Kubernetes scheduler to start the application.

D- Lift and shift to a VM on Compute Engine. Use an instance schedule to start and stop the instance.

Answer:

B

Question 9

Question Type: MultipleChoice

You have a Bigtable instance that consists of three nodes that store personally identifiable information (Pll) dat

a. You need to log all read or write operations, including any metadata or configuration reads of this database table, in your company's Security Information and Event Management (SIEM) system. What should you do?

Options:

A- * Navigate to Cloud Mentioning in the Google Cloud console, and create a custom monitoring job for the Bigtable instance to track all changes.
* Create an alert by using webhook endpoints. with the SIEM endpoint as a receiver

B- Navigate to the Audit Logs page in the Google Cloud console, and enable Data Read. Data Write and Admin Read logs for the Bigtable instance
* Create a Pub/Sub topic as a Cloud Logging sink destination, and add your SIEM as a subscriber to the topic.

C- * Install the Ops Agent on the Bigtable instance during configuration. K
* Create a service account with read permissions for the Bigtable instance.
* Create a custom Dataflow job with this service account to export logs to the company's SIEM system.

D- * Navigate to the Audit Logs page in the Google Cloud console, and enable Admin Write logs for the Biglable instance.
* Create a Cloud Functions instance to export logs from Cloud Logging to your SIEM.

Answer:

B

Question 10

Question Type: MultipleChoice

You used the gcloud container clusters command to create two Google Cloud Kubernetes (GKE) clusters prod-cluster and dev-cluster.

* prod-cluster is a standard cluster.

* dev-cluster is an auto-pilot duster.

When you run the Kubect1 get nodes command, you only see the nodes from prod-cluster Which commands should you run to check the node status for dev-cluster?

A.

```
gcloud container clusters get-credentials dev-cluster
kubectl get nodes
```

B.

```
gcloud container clusters update -generate-password dev-cluster
kubectl get nodes
```

C.

```
kubectl config set-context dev-cluster
kubectl cluster-info
```

D.

```
kubectl config set-credentials dev-cluster
kubectl cluster-info
```

Options:

A- Option A

B- Option B

C- Option C

D- Option D

Answer:

C