



Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer

NEW QUESTION 1

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?

- A. Implement a Cloud Run job to rotate all service account keys periodically in pj-s
- B. Enforce an org policy to deny service account key creation with an exception to pj-sa.
- C. Implement a Kubernetes Cronjob to rotate all service account keys periodically
- D. Disable attachment of service accounts to resources in all projects with an exception to pj-sa.
- E. Enforce an org policy constraint allowing the lifetime of service account keys to be 24 hours
- F. Enforce an org policy constraint denying service account key creation with an exception on pj-sa.
- G. Enforce a DENY org policy constraint over the lifetime of service account keys for 24 hours
- H. 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 to true or false, or use a condition to apply the constraint to specific service accounts. By setting this constraint to false for 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 as 86400 for one day. By setting this constraint to 86400 for 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.

References:

- 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

NEW QUESTION 2

Your coworker has helped you set up several configurations for gcloud. You've noticed that you're running commands against the wrong project. Being new to the company, you haven't yet memorized any of the projects. With the fewest steps possible, what's the fastest way to switch to the correct configuration?

- A. Run `gcloud configurations list` followed by `gcloud configurations activate .`
- B. Run `gcloud config list` followed by `gcloud config activate.`
- C. Run `gcloud config configurations list` followed by `gcloud config configurations activate.`
- D. Re-authenticate with the `gcloud auth login` command and select the correct configurations on login.

Answer: C

Explanation:

as `gcloud config configurations list` can help check for the existing configurations and `activate` can help switch to the configuration.

`gcloud config configurations list` lists existing named configurations

`gcloud config configurations activate` activates an existing named configuration

Obtains access credentials for your user account via a web-based authorization flow. When this command completes successfully, it sets the active account in the current configuration to the account specified. If no configuration exists, it creates a configuration named default.

NEW QUESTION 3

Your company has embraced a hybrid cloud strategy where some of the applications are deployed on Google Cloud. A Virtual Private Network (VPN) tunnel connects your Virtual Private Cloud (VPC) in Google Cloud with your company's on-premises network. Multiple applications in Google Cloud need to connect to an on-premises database server, and you want to avoid having to change the IP configuration in all of your applications when the IP of the database changes.

What should you do?

- A. Configure Cloud NAT for all subnets of your VPC to be used when egressing from the VM instances.
- B. Create a private zone on Cloud DNS, and configure the applications with the DNS name.
- C. Configure the IP of the database as custom metadata for each instance, and query the metadata server.
- D. Query the Compute Engine internal DNS from the applications to retrieve the IP of the database.

Answer: B

Explanation:

Forwarding zones Cloud DNS forwarding zones let you configure target name servers for specific private zones. Using a forwarding zone is one way to implement outbound DNS forwarding from your VPC network. A Cloud DNS forwarding zone is a special type of Cloud DNS private zone. Instead of creating records within the zone, you specify a set of forwarding targets. Each forwarding target is an IP address of a DNS server, located in your VPC network, or in an on-premises network connected to your VPC network by Cloud VPN or Cloud Interconnect.

<https://cloud.google.com/nat/docs/overview>

DNS configuration Your on-premises network must have DNS zones and records configured so that Google domain names resolve to the set of IP addresses for either `private.googleapis.com` or `restricted.googleapis.com`. You can create Cloud DNS managed private zones and use a Cloud DNS inbound server policy, or you can configure on-premises name servers. For example, you can use BIND or Microsoft Active Directory DNS.

<https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid#config-domain>

NEW QUESTION 4

Your learn wants to deploy a specific content management system (CMS) solution to Google Cloud. You need a quick and easy way to deploy and install the solution. What should you do?

- A. Search for the CMS solution in Google Cloud Marketplac
- B. Use gcloud CLI to deploy the solution.
- C. Search for the CMS solution in Google Cloud Marketplac
- D. Deploy the solution directly from Cloud Marketplace.
- E. Search for the CMS solution in Google Cloud Marketplac
- F. Use Terraform and the Cloud Marketplace ID to deploy the solution with the appropriate parameters.
- G. Use the installation guide of the CMS provide
- H. Perform the installation through your configuration management system.

Answer: B

NEW QUESTION 5

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

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

Answer: D

Explanation:

➤ The Datastore emulator provides local emulation of the production Datastore environment. You can use the emulator to develop and test your application locally. Ref: <https://cloud.google.com/datastore/docs/tools/datastore-emulator>

NEW QUESTION 6

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. 1.00.0.0/0
- B. 10.0.0.0/8
- C. 172.16.0.0/12
- D. 192.168.0.0/16

Answer: B

Explanation:

https://cloud.google.com/vpc/docs/vpc#manually_created_subnet_ip_ranges

NEW QUESTION 7

Your company has developed a new application that consists of multiple microservices. You want to deploy the application to Google Kubernetes Engine (GKE), and you want to ensure that the cluster can scale as more applications are deployed in the future. You want to avoid manual intervention when each new application is deployed. What should you do?

- A. Deploy the application on GKE, and add a HorizontalPodAutoscaler to the deployment.
- B. Deploy the application on GKE, and add a VerticalPodAutoscaler to the deployment.
- C. Create a GKE cluster with autoscaling enabled on the node pool
- D. Set a minimum and maximum for the size of the node pool.
- E. Create a separate node pool for each application, and deploy each application to its dedicated node pool.

Answer: C

Explanation:

https://cloud.google.com/kubernetes-engine/docs/how-to/cluster-autoscaler#adding_a_node_pool_with_autoscal

NEW QUESTION 8

Your team is using Linux instances on Google Cloud. You need to ensure that your team logs in to these instances in the most secure and cost efficient way. What should you do?

- A. Attach a public IP to the instances and allow incoming connections from the internet on port 22 for SSH.
- B. Use a third party tool to provide remote access to the instances.
- C. Use the `gcloud compute ssh` command with the `--tunnel-through-iap` flag
- D. Allow ingress traffic from the IP range 35.235.240.0/20 on port 22.
- E. Create a bastion host with public internet access
- F. Create the SSH tunnel to the instance through the bastion host.

Answer: D

NEW QUESTION 9

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

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

Answer: B

Explanation:

Coldline Storage is the perfect service to store audit logs from all the projects and is very cost-efficient as well. Coldline Storage is a very low-cost, highly durable storage service for storing infrequently accessed data.

NEW QUESTION 10

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
- B. Give the private key to each member of your team
- 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 key
- 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 account
- G. Grant the "compute.osAdminLogin" role to the Google group corresponding to this team.
- H. Generate a new SSH key pair
- I. Give the private key to each member of your team
- 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 10

Your organization has a dedicated person who creates and manages all service accounts for Google Cloud projects. You need to assign this person the minimum role for projects. What should you do?

- A. Add the user to roles/iam.roleAdmin role.
- B. Add the user to roles/iam.securityAdmin role.
- C. Add the user to roles/iam.serviceAccountUser role.
- D. Add the user to roles/iam.serviceAccountAdmin role.

Answer: D

NEW QUESTION 13

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

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

Answer: A

Explanation:

<https://cloud.google.com/iam/docs/understanding-roles#billing-roles>

NEW QUESTION 15

Your finance team wants to view the billing report for your projects. You want to make sure that the finance team does not get additional permissions to the project. What should you do?

- A. Add the group for the finance team to roles/billing user role.
- B. Add the group for the finance team to roles/billing admin role.
- C. Add the group for the finance team to roles/billing viewer role.
- D. Add the group for the finance team to roles/billing project/Manager role.

Answer: C

Explanation:

"Billing Account Viewer access would usually be granted to finance teams, it provides access to spend information, but does not confer the right to link or unlink projects or otherwise manage the properties of the billing account." <https://cloud.google.com/billing/docs/how-to/billing-access>

NEW QUESTION 16

Your company's security vulnerability management policy wants 3 member of the security team to have visibility into vulnerabilities and other OS metadata for a specific Compute Engine instance. This Compute Engine instance hosts a critical application in your Goggle Cloud project. You need to implement your company's security vulnerability management policy. What should you do?

- A. • Ensure that the Ops Agent is installed on the Compute Engine instance. • Create a custom metric in the Cloud Monitoring dashboard. • Provide the security team member with access to this dashboard.
- B. • Ensure that the Ops Agent is installed on the Compute Engine instance. • Provide the security team member roles/configure.inventoryViewer permission.
- C. • Ensure that the OS Config agent is installed on the Compute Engine instance. • Provide the security team member roles/configure.vulnerabilityViewer permission.
- D. • Ensure that the OS Config agent is installed on the Compute Engine instance. • Create a log sink to a BigQuery dataset. • Provide the security team member with access to this dataset.

Answer: C

NEW QUESTION 18

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?

- 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.

References:

- > 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>

NEW QUESTION 22

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: B

Explanation:

<https://cloud.google.com/iam/docs/best-practices-for-using-and-managing-service-accounts>

If an application uses third-party or custom identities and needs to access a resource, such as a BigQuery dataset or a Cloud Storage bucket, it must perform a transition between principals. Because Google Cloud APIs don't recognize third-party or custom identities, the application can't propagate the end-user's identity to BigQuery or Cloud Storage. Instead, the application has to perform the access by using a different Google identity.

NEW QUESTION 25

You are building a data lake on Google Cloud for your Internet of Things (IoT) application. The IoT application has millions of sensors that are constantly streaming structured and unstructured data to your backend in the cloud. You want to build a highly available and resilient architecture based on Google-recommended practices. What should you do?

- A. Stream data to Pub/Sub, and use Dataflow to send data to Cloud Storage.
- B. Stream data to Pub/Sub.
- C. and use Storage Transfer Service to send data to BigQuery.
- D. Stream data to Dataflow, and use Storage Transfer Service to send data to BigQuery.
- E. Stream data to Dataflow, and use Dataprep by Trifacta to send data to Bigtable.

Answer: B

NEW QUESTION 30

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

NEW QUESTION 32

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

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

Answer: A

Explanation:

Choosing a region and zone You choose which region or zone hosts your resources, which controls where your data is stored and used. Choosing a region and zone is important for several reasons:

Handling failures

Distribute your resources across multiple zones and regions to tolerate outages. Google designs zones to be independent from each other: a zone usually has power, cooling, networking, and control planes that are isolated from other zones, and most single failure events will affect only a single zone. Thus, if a zone becomes unavailable, you can transfer traffic to another zone in the same region to keep your services running. Similarly, if a region experiences any disturbances, you should have backup services running in a different region. For more information about distributing your resources and designing a robust system, see [Designing Robust Systems](#). Decreased network latency To decrease network latency, you might want to choose a region or zone that is close to your point of service.

https://cloud.google.com/compute/docs/regions-zones#choosing_a_region_and_zone

NEW QUESTION 37

You have files in a Cloud Storage bucket that you need to share with your suppliers. You want to restrict the time that the files are available to your suppliers to 1 hour. You want to follow Google recommended practices. What should you do?

- A. Create a service account with just the permissions to access files in the bucket.
- B. Create a JSON key for the service account.
- C. Execute the command `gsutil signurl -m 1h gs://**`.
- D. Create a service account with just the permissions to access files in the bucket.
- E. Create a JSON key for the service account.
- F. Execute the command `gsutil signurl -d 1h gs://**`.
- G. Create a service account with just the permissions to access files in the bucket.
- H. Create a JSON key for the service account.
- I. Execute the command `gsutil signurl -p 60m gs://`.
- J. Create a JSON key for the Default Compute Engine Service Account.
- K. Execute the command `gsutil signurl -t 60m gs://***`.

Answer: B

Explanation:

This command correctly specifies the duration that the signed url should be valid for by using the -d flag. The default is 1 hour so omitting the -d flag would have also resulted in the same outcome. Times may be specified with no suffix (default hours), or with s = seconds, m = minutes, h = hours, d = days. The max duration allowed is 7d. Ref: <https://cloud.google.com/storage/docs/gsutil/commands/signurl>

NEW QUESTION 38

A team of data scientists infrequently needs to use a Google Kubernetes Engine (GKE) cluster that you manage. They require GPUs for some long-running, non-restartable jobs. You want to minimize cost. What should you do?

- A. Enable node auto-provisioning on the GKE cluster.
- B. Create a VerticalPodAutoscaler for those workloads.
- C. Create a node pool with preemptible VMs and GPUs attached to those VMs.
- D. Create a node pool of instances with GPUs, and enable autoscaling on this node pool with a minimum size of 1.

Answer: A

Explanation:

auto-provisioning = Attaches and deletes node pools to cluster based on the requirements. Hence creating a GPU node pool, and auto-scaling would be better
<https://cloud.google.com/kubernetes-engine/docs/how-to/node-auto-provisioning>

NEW QUESTION 39

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 template
- 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 version
- G. Update the existing managed instance group with the new instance template
- 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

Explanation:

https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups#max_

NEW QUESTION 41

You installed the Google Cloud CLI on your workstation and set the proxy configuration. However, you are worried that your proxy credentials will be recorded in the gcloud CLI logs. You want to prevent your proxy credentials from being logged. What should you do?

- A. Configure username and password by using gcloud configure set proxy/username and gcloud configure set proxy/ proxy/password commands.
- B. Encode username and password in sha256 encoding, and save it to a text file
- C. Use filename as a value in the gcloud configure set core/custom_ca_certs_file command.
- D. Provide values for CLOUDSDK_USERNAME and CLOUDSDK_PASSWORD in the gcloud CLI tool configure file.
- E. Set the CLOUDSDK_PROXY_USERNAME and CLOUDSDK_PROXY_PASSWORD properties by using environment variables in your command line tool.

Answer: D

NEW QUESTION 43

The sales team has a project named Sales Data Digest that has the ID acme-data-digest. You need to set up similar Google Cloud resources for the marketing team but their resources must be organized independently of the sales team. What should you do?

- A. Grant the Project Editor role to the Marketing team for acme data digest
- B. Create a Project Lien on acme-data digest and then grant the Project Editor role to the Marketing team
- C. Create another project with the ID acme-marketing-data-digest for the Marketing team and deploy the resources there
- D. Create a new project named Meeting Data Digest and use the ID acme-data-digest. Grant the Project Editor role to the Marketing team.

Answer: C

NEW QUESTION 44

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 the open source CLI tool
- E. Snapshot Debugger, to find out why the data was not refreshed correctly.

Answer: D

Explanation:

Cloud Debugger helps inspect the state of an application, at any code location, without stopping or slowing down the running app // <https://cloud.google.com/debugger/docs>

NEW QUESTION 49

Your company has workloads running on Compute Engine and on-premises. The Google Cloud Virtual Private Cloud (VPC) is connected to your WAN over a Virtual Private Network (VPN). You need to deploy a new Compute Engine instance and ensure that no public Internet traffic can be routed to it. What should you do?

- A. Create the instance without a public IP address.
- B. Create the instance with Private Google Access enabled.
- C. Create a deny-all egress firewall rule on the VPC network.
- D. Create a route on the VPC to route all traffic to the instance over the VPN tunnel.

Answer: A

Explanation:

VMs cannot communicate over the internet without a public IP address. Private Google Access permits access to Google APIs and services in Google's production infrastructure.

<https://cloud.google.com/vpc/docs/private-google-access>

NEW QUESTION 52

You need to manage a Cloud Spanner Instance for best query performance. Your instance in production runs in a single Google Cloud region. You need to improve performance in the shortest amount of time. You want to follow Google best practices for service configuration. What should you do?

- A. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 45% If you exceed this threshold, add nodes lo your instance.
- B. Create an alert in Cloud Monitoring to alert when the percentage to high priority CPU utilization reaches 45% Use database query statistics to identify queries that result in high CPU usage, and then rewrite those queries to optimize their resource usage
- C. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 65% If you exceed this threshold, add nodes to your instance
- D. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 65%. Use database query statistics to identity queries that result in high CPU usage, and then rewrite those queries to optimize their resource usage.

Answer: B

Explanation:

<https://cloud.google.com/spanner/docs/cpu-utilization#recommended-max>

NEW QUESTION 54

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

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

Answer: B

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset> https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset#usage_patterns
DaemonSets attempt to adhere to a one-Pod-per-node model, either across the entire cluster or a subset of nodes. As you add nodes to a node pool, DaemonSets automatically add Pods to the new nodes as needed.

In GKE, DaemonSets manage groups of replicated Pods and adhere to a one-Pod-per-node model, either across the entire cluster or a subset of nodes. As you add nodes to a node pool, DaemonSets automatically add Pods to the new nodes as needed. So, this is a perfect fit for our monitoring pod.

Ref: <https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset>

DaemonSets are useful for deploying ongoing background tasks that you need to run on all or certain nodes, and which do not require user intervention. Examples of such tasks include storage daemons like ceph, log collection daemons like fluentd, and node monitoring daemons like collectd. For example, you could have DaemonSets for each type of daemon run on all of your nodes. Alternatively, you could run multiple DaemonSets for a single type of daemon, but have them use different configurations for different hardware types and resource needs.

NEW QUESTION 55

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

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

Answer: B

NEW QUESTION 58

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

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

Answer: D

NEW QUESTION 61

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:

Managed instance groups offer autoscaling capabilities that let you automatically add or delete instances from a managed instance group based on increases or decreases in load (CPU Utilization in this case). Autoscaling helps your apps gracefully handle increases in traffic and reduce costs when the need for resources is lower. You define the autoscaling policy and the autoscaler performs automatic scaling based on the measured load (CPU Utilization in this case). Autoscaling works by adding more instances to your instance group when there is more load (upscaling), and deleting instances when the need for instances is lowered (downscaling). Ref: <https://cloud.google.com/compute/docs/autoscaler>

NEW QUESTION 65

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

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

Answer: B

Explanation:

You can generate Windows passwords using either the Google Cloud Console or the `gcloud` command-line tool. This option uses the right syntax to reset the windows password.

`gcloud compute reset-windows-password windows-instance`

Ref: <https://cloud.google.com/compute/docs/instances/windows/creating-passwords-for-windows-instances#gc>

NEW QUESTION 68

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

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

Answer: A

Explanation:

<https://cloud.google.com/compute/docs/troubleshooting/troubleshooting-migs> https://cloud.google.com/compute/docs/instance-templates#how_to_update_instance_templates

NEW QUESTION 70

You are building an application that processes data files uploaded from thousands of suppliers. Your primary goals for the application are data security and the expiration of aged data. You need to design the application to:

- Restrict access so that suppliers can access only their own data.
- Give suppliers write access to data only for 30 minutes.
- Delete data that is over 45 days old.

You have a very short development cycle, and you need to make sure that the application requires minimal maintenance. Which two strategies should you use? (Choose two.)

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

Answer: AB

Explanation:

(A) Object Lifecycle Management Delete

The Delete action deletes an object when the object meets all conditions specified in the lifecycle rule.

Exception: In buckets with Object Versioning enabled, deleting the live version of an object causes it to become a noncurrent version, while deleting a noncurrent version deletes that version permanently.

<https://cloud.google.com/storage/docs/lifecycle#delete>

(B) Signed URLs

This page provides an overview of signed URLs, which you use 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>

NEW QUESTION 71

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

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

Answer: A

Explanation:

Change your default zone and region in the metadata server Note: This only applies to the default configuration. You can change the default zone and region in

your metadata server by making a request to the metadata server. For example: `gcloud compute project-info add-metadata \ --metadata google-compute-default-region=europe-west1,google-compute-default-zone=europe-west1-b` The gcloud command-line tool only picks up on new default zone and region changes after you rerun the gcloud init command. After updating your default metadata, run gcloud init to reinitialize your default configuration.
https://cloud.google.com/compute/docs/gcloud-compute#change_your_default_zone_and_region_in_the_metad

NEW QUESTION 72

Your customer has implemented a solution that uses Cloud Spanner and notices some read latency-related performance issues on one table. This table is accessed only by their users using a primary key. The table schema is shown below.

```
CREATE TABLE Persons (
  person_id INT64 NOT NULL,    // sequential number based on number of registration
  account_creation_date DATE, // system date
  birthdate DATE,            // customer birthdate
  firstname STRING (255),    // first name
  lastname STRING (255),     // last name
  profile_picture BYTES (255) // profile picture
) PRIMARY KEY (person_id)
```

You want to resolve the issue. What should you do?

- A. Remove the profile_picture field from the table.
- B. Add a secondary index on the person_id column.
- C. Change the primary key to not have monotonically increasing values.
- D. Create a secondary index using the following Data Definition Language (DDL):

```
CREATE INDEX person_id_ix
ON Persons (
  person_id,
  firstname,
  lastname
) STORING (
  profile_picture
)
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

Explanation:

As mentioned in Schema and data model, you should be careful when choosing a primary key to not accidentally create hotspots in your database. One cause of hotspots is having a column whose value monotonically increases as the first key part, because this results in all inserts occurring at the end of your key space. This pattern is undesirable because Cloud Spanner divides data among servers by key ranges, which means all your inserts will be directed at a single server that will end up doing all the work. <https://cloud.google.com/spanner/docs/schema-design#primary-key-prevent-hotspots>

NEW QUESTION 74

The core business of your company is to rent out construction equipment at a large scale. All the equipment that is being rented out has been equipped with multiple sensors that send event information every few seconds. These signals can vary from engine status, distance traveled, fuel level, and more. Customers are billed based on the consumption monitored by these sensors. You expect high throughput – up to thousands of events per hour per device – and need to retrieve consistent data based on the time of the event. Storing and retrieving individual signals should be atomic. What should you do?

- A. Create a file in Cloud Storage per device and append new data to that file.
- B. Create a file in Cloud Filestore per device and append new data to that file.
- C. Ingest the data into Datastor
- D. Store data in an entity group based on the device.
- E. Ingest the data into Cloud Bigtabl
- F. Create a row key based on the event timestamp.

Answer: D

Explanation:

Keyword need to look for

- "High Throughput",
- "Consistent",
- "Property based data insert/fetch like engine status, distance traveled, fuel level, and more." which can be designed in column,
- "Large Scale Customer Base + Each Customer has multiple sensor which send event in seconds" This will go for pera bytes situation,
- Export data based on the time of the event.
- Atomic
- o BigTable will fit all requirement.
- o DataStore is not fully Atomic
- o CloudStorage is not a option where we can export data based on time of event. We need another solution to do that
- o FireStore can be used with MobileSDK.

NEW QUESTION 79

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

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

Answer: B

Explanation:

Billing Administrators can not create a new billing account, and the project is presumably already created. Project Billing Manager allows you to link the created billing account to the project. It is vague on how the billing account gets created but by process of elimination

NEW QUESTION 83

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

Explanation:

IP 10.0.3.21 is internal by default, and to ensure that it will be static non-changing it should be selected as static internal ip address.

NEW QUESTION 84

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 the 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: B

Explanation:

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

NEW QUESTION 85

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 Organization
- B. Create a new VPC and add all instances.
- C. Verify that both projects are in a GCP Organization
- 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

Explanation:

Shared VPC allows an organization to connect resources from multiple projects to a common Virtual Private Cloud (VPC) network, so that they can communicate with each other securely and efficiently using internal IPs from that network. When you use Shared VPC, you designate a project as a host project and attach one or more other service projects to it. The VPC networks in the host project are called Shared VPC networks. Eligible resources from service projects can use subnets in the Shared VPC network

<https://cloud.google.com/vpc/docs/shared-vpc>

"For example, an existing instance in a service project cannot be reconfigured to use a Shared VPC network, but a new instance can be created to use available subnets in a Shared VPC network."

NEW QUESTION 89

Your company developed an application to deploy on Google Kubernetes Engine. Certain parts of the application are not fault-tolerant and are allowed to have downtime. Other parts of the application are critical and must always be available. You need to configure a Google Kubernetes Engine cluster while optimizing for cost. What should you do?

- A. Create a cluster with a single node-pool by using standard VM
- B. Label the fault-tolerant Deployments as spot-true.

- C. Create a cluster with a single node-pool by using Spot VM
- D. Label the critical Deployments as spot-false.
- E. Create a cluster with both a Spot VM node pool and a node pool by using standard VMs Deploy the critical
- F. deployments on the Spot VM node pool and the fault-tolerant deployments on the node pool by using standard VMs.
- G. Create a cluster with both a Spot VM node pool and by using standard VM
- H. Deploy the critical deployments on the node pool by using standard VMs and the fault-tolerant deployments on the Spot VM node pool.

Answer: C

NEW QUESTION 91

Your company completed the acquisition of a startup and is now merging the IT systems of both companies. The startup had a production Google Cloud project in their organization. You need to move this project into your organization and ensure that the project is billed to your organization. You want to accomplish this task with minimal effort. What should you do?

- A. Use the project
- B. move method to move the project to your organization
- C. Update the billing account of the project to that of your organization.
- D. Ensure that you have an Organization Administrator Identity and Access Management (IAM) role assigned to you in both organization
- E. Navigate to the Resource Manager in the startup's Google Cloud organization, and drag the project to your company's organization.
- F. Create a Private Catalog for the Google Cloud Marketplace, and upload the resources of the startup's production project to the Catalog
- G. Share the Catalog with your organization, and deploy the resources in your company's project.
- H. Create an infrastructure-as-code template for all resources in the project by using Terraform
- I. and deploy that template to a new project in your organization
- J. Delete the project from the startup's Google Cloud organization.

Answer: A

NEW QUESTION 92

You are configuring Cloud DNS. You want to create DNS records to point home.mydomain.com, mydomain.com. and www.mydomain.com to the IP address of your Google Cloud load balancer. What should you do?

- A. Create one CNAME record to point mydomain.com to the load balancer, and create two A records to point WWW and HOME to mydomain.com respectively.
- B. Create one CNAME record to point mydomain.com to the load balancer, and create two AAAA records to point WWW and HOME to mydomain.com respectively.
- C. Create one A record to point mydomain.com to the load balancer, and create two CNAME records to point WWW and HOME to mydomain.com respectively.
- D. Create one A record to point mydomain.com to the load balancer, and create two NS records to point WWW and HOME to mydomain.com respectively.

Answer: C

NEW QUESTION 97

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 is Google Cloud Memorystore?

Overview. Cloud Memorystore for Redis is a fully managed Redis service for Google Cloud Platform. Applications running on Google Cloud Platform can achieve extreme performance by leveraging the highly scalable, highly available, and secure Redis service without the burden of managing complex Redis deployments.

NEW QUESTION 101

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: C

Explanation:

From this source:

https://cloud.google.com/billing/docs/how-to/custom-roles#permission_association_and_inheritance

"For example, associating a project with a billing account requires the billing.resourceAssociations.create permission on the billing account and also the resourceManager.projects.createBillingAssignment permission on the project. This is because project permissions are required for actions where project owners control access, while billing account permissions are required for actions where billing account administrators control access. When both should be involved, both permissions are necessary."

NEW QUESTION 106

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 cluster.

When you run the `kubectl 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-credentials 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`

Answer: C

NEW QUESTION 109

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

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

Answer: A

Explanation:

performs a peering between the two VPC's (the statement makes sure that this option is feasible since it clearly specifies that there is no overlapping between the ip ranges of both vpc's), deploy the LoadBalancer as internal with the annotation, and configure the endpoint so that the compute engine instance can access the application internally, that is, without the need to have a public ip at any time and therefore, without the need to go outside the google network. The traffic, therefore, never crosses the public internet.

<https://medium.com/pablo-perez/k8s-externaltrafficpolicy-local-or-cluster-40b259a19404> <https://cloud.google.com/kubernetes-engine/docs/how-to/internal-load-balancing>

clients in a VPC network connected to the LoadBalancer network using VPC Network Peering can also access the Service <https://cloud.google.com/kubernetes-engine/docs/how-to/service-parameters>

NEW QUESTION 111

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

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

Answer: A

NEW QUESTION 115

You have been asked to set up Object Lifecycle Management for objects stored in storage buckets. The objects are written once and accessed frequently for 30 days. After 30 days, the objects are not read again unless there is a special need. The object should be kept for three years, and you need to minimize cost. What should you do?

- A. Set up a policy that uses Nearline storage for 30 days and then moves to Archive storage for three years.
- B. Set up a policy that uses Standard storage for 30 days and then moves to Archive storage for three years.
- C. Set up a policy that uses Nearline storage for 30 days, then moves to Coldline for one year, and then moves to Archive storage for two years.

D. Set up a policy that uses Standard storage for 30 days, then moves to Coldline for one year, and then moves to Archive storage for two years.

Answer: B

Explanation:

The key to understand the requirement is : "The objects are written once and accessed frequently for 30 days" Standard Storage
Standard Storage is best for data that is frequently accessed ("hot" data) and/or stored for only brief periods of time.

Archive Storage

Archive Storage is the lowest-cost, highly durable storage service for data archiving, online backup, and disaster recovery. Unlike the "coldest" storage services offered by other Cloud providers, your data is available within milliseconds, not hours or days. Archive Storage is the best choice for data that you plan to access less than once a year.

<https://cloud.google.com/storage/docs/storage-classes#standard>

NEW QUESTION 119

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

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

Answer: D

Explanation:

Install the workload in a compute engine VM, start and stop the instance as needed, because as per the question the VM runs for 30 hours, process can be performed offline and should not be interrupted, if interrupted we need to restart the batch process again. Preemptible VMs are cheaper, but they will not be available beyond 24hrs, and if the process gets interrupted the preemptible VM will restart.

NEW QUESTION 121

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

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

Answer: B

Explanation:

The resource that you need to get access is in the other project. roles/bigquery.dataViewer BigQuery Data Viewer

When applied to a table or view, this role provides permissions to: Read data and metadata from the table or view.

This role cannot be applied to individual models or routines. When applied to a dataset, this role provides permissions to:

Read the dataset's metadata and list tables in the dataset. Read data and metadata from the dataset's tables.

When applied at the project or organization level, this role can also enumerate all datasets in the project. Additional roles, however, are necessary to allow the running of jobs.

NEW QUESTION 124

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

- A. 1. Give the BigQuery Data Editor role on the platform-logs dataset to the service accounts used by your instances.2. Update your instances' metadata to add the following value: logs-destination:bq://platform-logs.
- B. 1. In Stackdriver Logging, create a logs export with a Cloud Pub/Sub topic called logs as a sink.2.Create a Cloud Function that is triggered by messages in the logs topic.3. Configure that Cloud Function to drop logs that are not from Compute Engine and to insert Compute Engine logs in the platform-logs dataset.
- C. 1. In Stackdriver Logging, create a filter to view only Compute Engine logs.2. Click Create Export.3.Choose BigQuery as Sink Service, and the platform-logs dataset as Sink Destination.
- D. 1. Create a Cloud Function that has the BigQuery User role on the platform-logs dataset.2. Configure this Cloud Function to create a BigQuery Job that executes this query:INSERT INTOdataset.platform-logs (timestamp, log)SELECT timestamp, log FROM compute.logsWHERE timestamp> DATE_SUB(CURRENT_DATE(), INTERVAL 1 DAY)3. Use Cloud Scheduler to trigger this Cloud Function once a day.

Answer: C

Explanation:

* 1. In Stackdriver Logging, create a filter to view only Compute Engine logs. 2. Click Create Export. 3. Choose BigQuery as Sink Service, and the platform-logs dataset as Sink Destination.

NEW QUESTION 126

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

Explanation:

"...we recommend that you enable Cloud Billing data export to BigQuery at the same time that you create a Cloud Billing account. "

<https://cloud.google.com/billing/docs/how-to/export-data-bigquery>

<https://medium.com/google-cloud/analyzing-google-cloud-billing-data-with-big-query-30bae1c2aae4>

NEW QUESTION 130

You are developing a financial trading application that will be used globally. Data is stored and queried using a relational structure, and clients from all over the world should get the exact identical state of the data. The application will be deployed in multiple regions to provide the lowest latency to end users. You need to select a storage option for the application data while minimizing latency. What should you do?

- A. Use Cloud Bigtable for data storage.
- B. Use Cloud SQL for data storage.
- C. Use Cloud Spanner for data storage.
- D. Use Firestore for data storage.

Answer: C

Explanation:

Keywords, Financial data (large data) used globally, data stored and queried using relational structure (SQL), clients should get exact identical copies(Strong Consistency), Multiple region, low latency to end user, select storage option to minimize latency.

NEW QUESTION 132

Your company's infrastructure is on-premises, but all machines are running at maximum capacity. You want to burst to Google Cloud. The workloads on Google Cloud must be able to directly communicate to the workloads on-premises using a private IP range. What should you do?

- A. In Google Cloud, configure the VPC as a host for Shared VPC.
- B. In Google Cloud, configure the VPC for VPC Network Peering.
- C. Create bastion hosts both in your on-premises environment and on Google Cloud.
- D. Configure both as proxy servers using their public IP addresses.
- E. Set up Cloud VPN between the infrastructure on-premises and Google Cloud.

Answer: D

Explanation:

"Google Cloud VPC Network Peering allows internal IP address connectivity across two Virtual Private Cloud (VPC) networks regardless of whether they belong to the same project or the same organization."

<https://cloud.google.com/vpc/docs/vpc-peering> while

"Cloud Interconnect provides low latency, high availability connections that enable you to reliably transfer data between your on-premises and Google Cloud Virtual Private Cloud (VPC) networks."

<https://cloud.google.com/network-connectivity/docs/interconnect/concepts/overview> and

"HA VPN is a high-availability (HA) Cloud VPN solution that lets you securely connect your on-premises network to your VPC network through an IPsec VPN connection in a single region."

<https://cloud.google.com/network-connectivity/docs/vpn/concepts/overview>

NEW QUESTION 136

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.databaseUser`
- B. Add the users to the role.
- C. Run `gcloud iam roles describe roles/spanner.databaseUser`
- D. Add the users to a new group
- 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: B

Explanation:

<https://cloud.google.com/spanner/docs/iam#spanner.databaseUser>

Using the `gcloud` tool, execute the `gcloud iam roles describe roles/spanner.databaseUser` command on Cloud Shell. Attach the users to a newly created Google group and add the group to the role.

NEW QUESTION 138

You are running a web application on Cloud Run for a few hundred users. Some of your users complain that the initial web page of the application takes much longer to load than the following pages. You want to follow Google's recommendations to mitigate the issue. What should you do?

- A. Update your web application to use the protocol HTTP/2 instead of HTTP/1.1
- B. Set the concurrency number to 1 for your Cloud Run service.
- C. Set the maximum number of instances for your Cloud Run service to 100.
- D. Set the minimum number of instances for your Cloud Run service to 3.

Answer: D

NEW QUESTION 139

You are managing several Google Cloud Platform (GCP) projects and need access to all logs for the past 60 days. You want to be able to explore and quickly analyze the log contents. You want to follow Google- recommended practices to obtain the combined logs for all projects. What should you do?

- A. Navigate to Stackdriver Logging and select resource.labels.project_id=""
- B. Create a Stackdriver Logging Export with a Sink destination to a BigQuery dataset
- C. Configure the table expiration to 60 days.
- D. Create a Stackdriver Logging Export with a Sink destination to Cloud Storage
- E. Create a lifecycle rule to delete objects after 60 days.
- F. Configure a Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery
- G. Configure the table expiration to 60 days.

Answer: B

Explanation:

➤ Navigate to Stackdriver Logging and select resource.labels.project_id=*. is not right.
 Log entries are held in Stackdriver Logging for a limited time known as the retention period which is 30 days (default configuration). After that, the entries are deleted. To keep log entries longer, you need to export them outside of Stackdriver Logging by configuring log sinks.
 Ref: <https://cloud.google.com/blog/products/gcp/best-practices-for-working-with-google-cloud-audit-logging>

➤ Configure a Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery. Configure the table expiration to 60 days. is not right.
 While this works, it makes no sense to use Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery when Google provides a feature (export sinks) that does exactly the same thing and works out of the box. Ref: https://cloud.google.com/logging/docs/export/configure_export_v2

➤ Create a Stackdriver Logging Export with a Sink destination to Cloud Storage. Create a lifecycle rule to delete objects after 60 days. is not right.
 You can export logs by creating one or more sinks that include a logs query and an export destination. Supported destinations for exported log entries are Cloud Storage, BigQuery, and Pub/Sub. Ref: https://cloud.google.com/logging/docs/export/configure_export_v2
 Sinks are limited to exporting log entries from the exact resource in which the sink was created: a Google Cloud project, organization, folder, or billing account. If it makes it easier to exporting from all projects of an organization, you can create an aggregated sink that can export log entries from all the projects, folders, and billing accounts of a Google Cloud organization. Ref: https://cloud.google.com/logging/docs/export/aggregated_sinks
 Either way, we now have the data in Cloud Storage, but querying logs information from Cloud Storage is harder than Querying information from BigQuery dataset. For this reason, we should prefer Big Query over Cloud Storage.

➤ Create a Stackdriver Logging Export with a Sink destination to a BigQuery dataset. Configure the table expiration to 60 days. is the right answer.
 You can export logs by creating one or more sinks that include a logs query and an export destination. Supported destinations for exported log entries are Cloud Storage, BigQuery, and Pub/Sub. Ref: https://cloud.google.com/logging/docs/export/configure_export_v2
 Sinks are limited to exporting log entries from the exact resource in which the sink was created: a Google Cloud project, organization, folder, or billing account. If it makes it easier to exporting from all projects of an organization, you can create an aggregated sink that can export log entries from all the projects, folders, and billing accounts of a Google Cloud organization. Ref: https://cloud.google.com/logging/docs/export/aggregated_sinks
 Either way, we now have the data in a BigQuery Dataset. Querying information from a Big Query dataset is easier and quicker than analyzing contents in Cloud Storage bucket. As our requirement is to Quickly analyze the log contents, we should prefer Big Query over Cloud Storage.
 Also, You can control storage costs and optimize storage usage by setting the default table expiration for newly created tables in a dataset. If you set the property when the dataset is created, any table created in the dataset is deleted after the expiration period. If you set the property after the dataset is created, only new tables are deleted after the expiration period. For example, if you set the default table expiration to 7 days, older data is automatically deleted after 1 week. Ref: <https://cloud.google.com/bigquery/docs/best-practices-storage>

NEW QUESTION 143

You want to configure 10 Compute Engine instances for availability when maintenance occurs. Your requirements state that these instances should attempt to automatically restart if they crash. Also, the instances should be highly available including during system maintenance. What should you do?

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

Answer: A

Explanation:

Create an instance template for the instances so VMs have same specs. Set the "Automatic Restart" to on to VM automatically restarts upon crash. Set the "On-host maintenance" to Migrate VM instance. This will take care of VM during maintenance window. It will migrate VM instance making it highly available Add the instance template to an instance group so instances can be managed.

- onHostMaintenance: Determines the behavior when a maintenance event occurs that might cause your instance to reboot.
- [Default] MIGRATE, which causes Compute Engine to live migrate an instance when there is a maintenance event.
- TERMINATE, which stops an instance instead of migrating it.
- automaticRestart: Determines the behavior when an instance crashes or is stopped by the system.
- [Default] true, so Compute Engine restarts an instance if the instance crashes or is stopped.
- false, so Compute Engine does not restart an instance if the instance crashes or is stopped.

Enabling automatic restart ensures that compute engine instances are automatically restarted when they crash. And Enabling Migrate VM Instance enables live migrates i.e. compute instances are migrated during system maintenance and remain running during the migration.
 Automatic Restart If your instance is set to terminate when there is a maintenance event, or if your instance crashes because of an underlying hardware issue, you can set up Compute Engine to automatically restart the instance by setting the automaticRestart field to true. This setting does not apply if the instance is taken

offline through a user action, such as calling sudo shutdown, or during a zone outage. Ref: <https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options#autorestart>

Enabling the Migrate VM Instance option migrates your instance away from an infrastructure maintenance event, and your instance remains running during the migration. Your instance might experience a short period of decreased performance, although generally, most instances should not notice any difference. This is ideal for instances that require constant uptime and can tolerate a short period of decreased performance. Ref: https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options#live_

NEW QUESTION 148

You are designing an application that uses WebSockets and HTTP sessions that are not distributed across the web servers. You want to ensure the application runs properly on Google Cloud Platform. What should you do?

- A. Meet with the cloud enablement team to discuss load balancer options.
- B. Redesign the application to use a distributed user session service that does not rely on WebSockets and HTTP sessions.
- C. Review the encryption requirements for WebSocket connections with the security team.
- D. Convert the WebSocket code to use HTTP streaming.

Answer: A

Explanation:

> Google HTTP(S) Load Balancing has native support for the WebSocket protocol when you use HTTP or HTTPS, not HTTP/2, as the protocol to the backend. Ref: https://cloud.google.com/load-balancing/docs/https#websocket_proxy_support

> We dont need to convert WebSocket code to use HTTP streaming or Redesign the application, as WebSocket support is offered by Google HTTP(S) Load Balancing. Reviewing the encryption requirements is a good idea but it has nothing to do with WebSockets.

NEW QUESTION 149

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

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

Answer: A

Explanation:

<https://cloud.google.com/storage/docs/audit-logs> https://cloud.google.com/compute/docs/logging/audit-logging#audited_operations

NEW QUESTION 154

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 company.
- E. Once they grant you that role, replicate the setup you have in the development environment in that project.

Answer: A

Explanation:

This aligns with Googles recommended practices. By creating a new project, we achieve complete isolation between development and production environments; as well as isolate this production application from production applications of other departments.

Ref: <https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#define-hierarchy>

NEW QUESTION 157

You have an application that uses Cloud Spanner as a backend database. The application has a very predictable traffic pattern. You want to automatically scale up or down the number of Spanner nodes depending on traffic. What should you do?

- A. Create a cron job that runs on a scheduled basis to review stackdriver monitoring metrics, and then resize the Spanner instance accordingly.
- B. Create a Stackdriver alerting policy to send an alert to oncall SRE emails when Cloud Spanner CPU exceeds the threshold
- C. SREs would scale resources up or down accordingly.
- D. Create a Stackdriver alerting policy to send an alert to Google Cloud Support email when Cloud Spanner CPU exceeds your threshold
- E. Google support would scale resources up or down accordingly.
- F. Create a Stackdriver alerting policy to send an alert to webhook when Cloud Spanner CPU is over or under your threshold
- G. Create a Cloud Function that listens to HTTP and resizes Spanner resources accordingly.

Answer: D

Explanation:

As to mexblood1's point, CPU utilization is a recommended proxy for traffic when it comes to Cloud Spanner. See: Alerts for high CPU utilization The following

table specifies our recommendations for maximum CPU usage for both single-region and multi-region instances. These numbers are to ensure that your instance has enough compute capacity to continue to serve your traffic in the event of the loss of an entire zone (for single-region instances) or an entire region (for multi-region instances). - <https://cloud.google.com/spanner/docs/cpu-utilization>

NEW QUESTION 158

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

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

Answer: D

Explanation:

<https://cloud.google.com/compute/docs/instances/connecting-to-windows#remote-desktop-connection-app>

<https://cloud.google.com/compute/docs/instances/windows/generating-credentials> <https://cloud.google.com/compute/docs/instances/connecting-to-windows#before-you-begin>

NEW QUESTION 159

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

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

Answer: A

NEW QUESTION 160

Your company wants to migrate their on-premises workloads to Google Cloud. The current on-premises workloads consist of:

- A Flask web application
- AbackendAPI
- A scheduled long-running background job for ETL and reporting.

You need to keep operational costs low You want to follow Google-recommended practices to migrate these workloads to serverless solutions on Google Cloud. What should you do?

- A. Migrate the web application to App Engine and the backend API to Cloud Run Use Cloud Tasks to run your background job on Compute Engine
- B. Migrate the web application to App Engine and the backend API to Cloud Ru
- C. Use Cloud Tasks to run your background job on Cloud Run.
- D. Run the web application on a Cloud Storage bucket and the backend API on Cloud Run Use Cloud Tasks to run your background job on Cloud Run.
- E. Run the web application on a Cloud Storage bucket and the backend API on Cloud Ru
- F. Use Cloud Tasks to run your background job on Compute Engine

Answer: B

NEW QUESTION 164

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:

<https://cloud.google.com/deployment-manager/docs/configuration/create-basic-configuration>

NEW QUESTION 165

You need to host an application on a Compute Engine instance in a project shared with other teams. You want to prevent the other teams from accidentally causing downtime on that application. Which feature should you use?

- A. Use a Shielded VM.
- B. Use a Preemptible VM.

- C. Use a sole-tenant node.
- D. Enable deletion protection on the instance.

Answer: D

Explanation:

As part of your workload, there might be certain VM instances that are critical to running your application or services, such as an instance running a SQL server, a server used as a license manager, and so on. These VM instances might need to stay running indefinitely so you need a way to protect these VMs from being deleted. By setting the deletionProtection flag, a VM instance can be protected from accidental deletion. If a user attempts to delete a VM instance for which you have set the deletionProtection flag, the request fails. Only a user that has been granted a role with compute.instances.create permission can reset the flag to allow the resource to be deleted. Ref: <https://cloud.google.com/compute/docs/instances/preventing-accidental-vm-deletion>

NEW QUESTION 169

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 cluster
- F. Ask your ML team to add the cloud.google.com/gke -accelerator: nvidia-tesla-p100 nodeSelector to their pod specification.

Answer: D

Explanation:

This is the most optimal solution. Rather than recreating all nodes, you create a new node pool with GPU enabled. You then modify the pod specification to target particular GPU types by adding node selector to your workloads Pod specification. You still have a single cluster so you pay Kubernetes cluster management fee for just one cluster thus minimizing the cost. Ref: <https://cloud.google.com/kubernetes-engine/docs/how-to/gpus> Ref: <https://cloud.google.com/kubern>

Example:

```
> apiVersion: v1
> kind: Pod
> metadata:
> name: my-gpu-pod
> spec:
> containers:
> name: my-gpu-container
> image: nvidia/cuda:10.0-runtime-ubuntu18.04
> command: [/bin/bash]
> resources:
> limits:
> nvidia.com/gpu: 2
> nodeSelector:
> cloud.google.com/gke-accelerator: nvidia-tesla-k80 # or nvidia-tesla-p100 or nvidia-tesla-p4 or nvidia-tesla-v100 or nvidia-tesla-t4
```

NEW QUESTION 174

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:

(https://cloud.google.com/resource-manager/docs/project-migration#change_billing_account) <https://cloud.google.com/billing/docs/concepts>
<https://cloud.google.com/resource-manager/docs/project-migration>

NEW QUESTION 175

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: B

Explanation:

Cloud Spanner is a relational database and is highly scalable. Cloud Spanner is a highly scalable,

enterprise-grade, globally-distributed, and strongly consistent database service built for the cloud specifically to combine the benefits of relational database structure with a non-relational horizontal scale. This combination delivers high-performance transactions and strong consistency across rows, regions, and continents with an industry-leading 99.999% availability SLA, no planned downtime, and enterprise-grade security

Ref: <https://cloud.google.com/spanner>

Graphical user interface, application, Teams Description automatically generated

	CLOUD SPANNER	TRADITIONAL RELATIONAL	TRADITIONAL NON-RELATIONAL
Schema	✓ Yes	✓ Yes	✗ No
SQL	✓ Yes	✓ Yes	✗ No
Consistency	✓ Strong	✓ Strong	✗ Eventual
Availability	✓ High	✗ Failover	✓ High
Scalability	✓ Horizontal	✗ Vertical	✓ Horizontal
Replication	✓ Automatic	⚙️ Configurable	⚙️ Configurable

NEW QUESTION 180

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

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

Answer: A

Explanation:

https://cloud.google.com/compute/docs/instance-groups/autohealing-instances-in-migs#setting_up_an_autoheali

NEW QUESTION 184

You want to select and configure a cost-effective solution for relational data on Google Cloud Platform. You are working with a small set of operational data in one geographic location. You need to support point-in-time recovery. What should you do?

- A. Select Cloud SQL (MySQL). Verify that the enable binary logging option is selected.
- B. Select Cloud SQL (MySQL). Select the create failover replicas option.
- C. Select Cloud Spanne
- D. Set up your instance with 2 nodes.
- E. Select Cloud Spanne
- F. Set up your instance as multi-regional.

Answer: A

NEW QUESTION 185

You want to permanently delete a Pub/Sub topic managed by Config Connector in your Google Cloud project. What should you do?

- A. Use kubectl to delete the topic resource.
- B. Use gcloud CLI to delete the topic.
- C. Use kubectl to create the label deleted-by-cnrm and to change its value to true for the topic resource.
- D. Use gcloud CLI to update the topic label managed-by-cnrm to false.

Answer: A

NEW QUESTION 187

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 192

You have a website hosted on App Engine standard environment. You want 1% of your users to see a new test version of the website. You want to minimize complexity. What should you do?

- A. Deploy the new version in the same application and use the `--migrate` option.
- B. Deploy the new version in the same application and use the `--splits` option to give a weight of 99 to the current version and a weight of 1 to the new version.
- C. Create a new App Engine application in the same projec

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

Answer: B

Explanation:

<https://cloud.google.com/appengine/docs/standard/python/splitting-traffic#gcloud>

NEW QUESTION 197

You need to track and verify modifications to a set of Google Compute Engine instances in your Google Cloud project. In particular, you want to verify OS system patching events on your virtual machines (VMs). What should you do?

- A. Review the Compute Engine activity logs. Select and review the Admin Event logs.
- B. Review the Compute Engine activity logs. Select and review the System Event logs.
- C. Install the Cloud Logging Agent. In Cloud Logging, review the Compute Engine syslog logs.
- D. Install the Cloud Logging Agent. In Cloud Logging, review the Compute Engine operation logs.

Answer: A

NEW QUESTION 202

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 to 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 topic.
- E. Create an application that sends all medical images to the Pub/Sub topic.

Answer: C

Explanation:

They require cloud storage for archival and they want to automate the process to upload new medical images to cloud storage, hence we go for gsutil to copy on-prem images to cloud storage and automate the process via cron job. Whereas Pub/Sub listens to the changes in the Cloud Storage bucket and triggers the pub/sub topic, which is not required.

NEW QUESTION 206

An application generates daily reports in a Compute Engine virtual machine (VM). The VM is in the project corp-iot-insights. Your team operates only in the project corp-aggregate-reports and needs a copy of the daily exports in the bucket corp-aggregate-reports-storage. You want to configure access so that the daily reports from the VM are available in the bucket corp-aggregate-reports-storage and use as few steps as possible while following Google-recommended practices. What should you do?

- A. Move both projects under the same folder.
- B. Grant the VM Service Account the role Storage Object Creator on corp-aggregate-reports-storage.
- C. Create a Shared VPC network between both projects.
- D. Grant the VM Service Account the role Storage Object Creator on corp-iot-insights.
- E. Make corp-aggregate-reports-storage public and create a folder with a pseudo-randomized suffix name. Share the folder with the IoT team.

Answer: B

Explanation:

Predefined roles

The following table describes Identity and Access Management (IAM) roles that are associated with Cloud Storage and lists the permissions that are contained in each role. Unless otherwise noted, these roles can be applied either to entire projects or specific buckets.

Storage Object Creator (roles/storage.objectCreator) Allows users to create objects. Does not give permission to view, delete, or overwrite objects.

<https://cloud.google.com/storage/docs/access-control/iam-roles#standard-roles>

NEW QUESTION 210

An external member of your team needs list access to compute images and disks in one of your projects. You want to follow Google-recommended practices when you grant the required permissions to this user. What should you do?

- A. Create a custom role, and add all the required compute.disks.list and compute.images.list permissions as includedPermissions.
- B. Grant the custom role to the user at the project level.
- C. Create a custom role based on the Compute Image User role. Add the compute.disks.list to the includedPermissions field. Grant the custom role to the user at the project level.
- D. Grant the Compute Storage Admin role at the project level.
- E. Create a custom role based on the Compute Storage Admin role.
- F. Exclude unnecessary permissions from the custom role.
- G. Grant the custom role to the user at the project level.

Answer: B

NEW QUESTION 215

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

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

Answer: A

Explanation:

<https://cloud.google.com/sdk/gcloud> https://cloud.google.com/sdk/docs/configurations#multiple_configurations

NEW QUESTION 217

You have a number of compute instances belonging to an unmanaged instances group. You need to SSH to one of the Compute Engine instances to run an ad hoc script. You've already authenticated gcloud, however, you don't have an SSH key deployed yet. In the fewest steps possible, what's the easiest way to SSH to the instance?

- A. Run gcloud compute instances list to get the IP address of the instance, then use the ssh command.
- B. Use the gcloud compute ssh command.
- C. Create a key with the ssh-keygen command
- D. Then use the gcloud compute ssh command.
- E. Create a key with the ssh-keygen command
- F. Upload the key to the instance
- G. Run gcloud compute instances list to get the IP address of the instance, then use the ssh command.

Answer: B

Explanation:

gcloud compute ssh ensures that the user's public SSH key is present in the project's metadata. If the user does not have a public SSH key, one is generated using ssh-keygen and added to the project's metadata. This is similar to the other option where we copy the key explicitly to the project's metadata but here it is done automatically for us. There are also security benefits with this approach. When we use gcloud compute ssh to connect to Linux instances, we are adding a layer of security by storing your host keys as guest attributes. Storing SSH host keys as guest attributes improve the security of your connections by helping to protect against vulnerabilities such as man-in-the-middle (MITM) attacks. On the initial boot of a VM instance, if guest attributes are enabled, Compute Engine stores your generated host keys as guest attributes.

Compute Engine then uses these host keys that were stored during the initial boot to verify all subsequent connections to the VM instance.

Ref: <https://cloud.google.com/compute/docs/instances/connecting-to-instance> Ref: <https://cloud.google.com/s>

NEW QUESTION 220

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

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

Answer: D

Explanation:

<https://gtseres.medium.com/using-service-accounts-across-projects-in-gcp-cf9473fef8f0#:~:text=Go%20to%20t>

NEW QUESTION 225

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?

- 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.

NEW QUESTION 228

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: D

Explanation:

https://cloud.google.com/bigquery/docs/access-control#custom_roles

Custom roles enable you to enforce the principle of least privilege, ensuring that the user and service accounts in your organization have only the permissions essential to performing their intended functions.

NEW QUESTION 233

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

Explanation:

Set budgets and budget alerts Overview Avoid surprises on your bill by creating Cloud Billing budgets to monitor all of your Google Cloud charges in one place. A budget enables you to track your actual Google Cloud spend against your planned spend. After you've set a budget amount, you set budget alert threshold rules that are used to trigger email notifications. Budget alert emails help you stay informed about how your spend is tracking against your budget. 2. Set budget scope Set the budget Scope and then click Next. In the Projects field, select one or more projects that you want to apply the budget alert to. To apply the budget alert to all the projects in the Cloud Billing account, choose Select all.

<https://cloud.google.com/billing/docs/how-to/budgets#budget-scop>

NEW QUESTION 235

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

NEW QUESTION 240

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

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

Answer: C

Explanation:

Using primitive roles The following table lists the primitive roles that you can grant to access a project, the description of what the role does, and the permissions bundled within that role. Avoid using primitive roles except when absolutely necessary. These roles are very powerful, and include a large number of permissions across all Google Cloud services. For more details on when you should use primitive roles, see the Identity and Access Management FAQ. IAM predefined roles are much more granular, and allow you to carefully manage the set of permissions that your users have access to. See Understanding Roles for a list of roles that can be granted at the project level. Creating custom roles can further increase the control you have over user permissions. https://cloud.google.com/resource-manager/docs/access-control-proj#using_primitive_roles

<https://cloud.google.com/iam/docs/understanding-custom-roles>

NEW QUESTION 242

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

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

Answer: B

Explanation:

Google Cloud Storage Triggers

Cloud Functions can respond to change notifications emerging from Google Cloud Storage. These notifications can be configured to trigger in response to various events inside a bucket—object creation, deletion, archiving and metadata updates.

Note: Cloud Functions can only be triggered by Cloud Storage buckets in the same Google Cloud Platform project.

Event types

Cloud Storage events used by Cloud Functions are based on Cloud Pub/Sub Notifications for Google Cloud Storage and can be configured in a similar way.

Supported trigger type values are: google.storage.object.finalize google.storage.object.delete google.storage.object.archive google.storage.object.metadataUpdate
 Object Finalize

Trigger type value: google.storage.object.finalize

This event is sent when a new object is created (or an existing object is overwritten, and a new generation of that object is created) in the bucket.

https://cloud.google.com/functions/docs/calling/storage#event_types

NEW QUESTION 247

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 and Delete 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

Explanation:

<https://cloud.google.com/storage/docs/lifecycle#setstorageclass-cost>

The object's time spent set at the original storage class counts towards any minimum storage duration that applies for the new storage class.

NEW QUESTION 248

You are migrating a business critical application from your local data center into Google Cloud. As part of your high-availability strategy, you want to ensure that any data used by the application will be immediately available if a zonal failure occurs. What should you do?

- A. Store the application data on a zonal persistent dis
- B. Create a snapshot schedule for the dis
- C. If an outage occurs, create a new disk from the most recent snapshot and attach it to a new VM in another zone.
- D. Store the application data on a zonal persistent dis
- E. If an outage occurs, create an instance in another zone with this disk attached.
- F. Store the application data on a regional persistent dis
- G. Create a snapshot schedule for the dis
- H. If an outage occurs, create a new disk from the most recent snapshot and attach it to a new VM in another zone.
- I. Store the application data on a regional persistent disk If an outage occurs, create an instance in another zone with this disk attached.

Answer: A

NEW QUESTION 250

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 grou
- D. Set the Autohealing health check to healthy (HTTP).
- E. Create a managed instance grou
- F. Verify that the autoscaling setting is on.

Answer: C

Explanation:

<https://cloud.google.com/compute/docs/instance-groups>

<https://cloud.google.com/load-balancing/docs/network/transition-to-backend-services#console>

➤ In order to enable auto-healing, you need to group the instances into a managed instance group.

Managed instance groups (MIGs) maintain the high availability of your applications by proactively keeping your virtual machine (VM) instances available. An auto-healing policy on the MIG relies on an application-based health check to verify that an application is responding as expected. If the auto-healer determines that an application isnt responding, the managed instance group automatically recreates that instance.

It is important to use separate health checks for load balancing and for auto-healing. Health checks for load balancing can and should be more aggressive because these health checks determine whether an instance receives user traffic. You want to catch non-responsive instances quickly, so you can redirect traffic if necessary. In contrast, health checking for auto-healing causes Compute Engine to proactively replace failing instances, so this health check should be more conservative than a load balancing health check.

NEW QUESTION 255

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. 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. Set the default key value project labels to department marketing for all services in this protect

Answer: A

NEW QUESTION 257

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

Associate-Cloud-Engineer Practice Exam Features:

- * Associate-Cloud-Engineer Questions and Answers Updated Frequently
- * Associate-Cloud-Engineer Practice Questions Verified by Expert Senior Certified Staff
- * Associate-Cloud-Engineer Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * Associate-Cloud-Engineer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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