

Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer



NEW QUESTION 1

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 2

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 3

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 4

You have one project called proj-sa where you manage all your service accounts. You want to be able to use a service account from this project to take snapshots of VMs running in another project called proj-vm. What should you do?

- A. Download the private key from the service account, and add it to each VMs custom metadata.
- B. Download the private key from the service account, and add the private key to each VM's SSH keys.
- C. Grant the service account the IAM Role of Compute Storage Admin in the project called proj-vm.
- D. When creating the VMs, set the service account's API scope for Compute Engine to read/write.

Answer: C

Explanation:

<https://gtseres.medium.com/using-service-accounts-across-projects-in-gcp-cf9473fef8f0>

You create the service account in proj-sa and take note of the service account email, then you go to proj-vm in IAM > ADD and add the service account's email as new member and give it the Compute Storage Admin role.

<https://cloud.google.com/compute/docs/access/iam#compute.storageAdmin>

NEW QUESTION 5

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 6

Your company has multiple projects linked to a single billing account in Google Cloud. You need to visualize the costs with specific metrics that should be dynamically calculated based on company-specific criteria. You want to automate the process. What should you do?

- A. In the Google Cloud console, visualize the costs related to the projects in the Reports section.
- B. In the Google Cloud console, visualize the costs related to the projects in the Cost breakdown section.
- C. In the Google Cloud console, use the export functionality of the Cost tab
- D. Create a Looker Studio dashboard on top of the CSV export.
- E. Configure Cloud Billing data export to BigQuery for the billing account
- F. Create a Looker Studio dashboard on top of the BigQuery export.

Answer: D

NEW QUESTION 7

You have an application on a general-purpose Compute Engine instance that is experiencing excessive disk read throttling on its Zonal SSD Persistent Disk. The application primarily reads large files from disk. The disk size is currently 350 GB. You want to provide the maximum amount of throughput while minimizing costs. What should you do?

- A. Increase the size of the disk to 1 TB.
- B. Increase the allocated CPU to the instance.
- C. Migrate to use a Local SSD on the instance.
- D. Migrate to use a Regional SSD on the instance.

Answer: C

Explanation:

Standard persistent disks are efficient and economical for handling sequential read/write operations, but they aren't optimized to handle high rates of random input/output operations per second (IOPS). If your apps require high rates of random IOPS, use SSD persistent disks. SSD persistent disks are designed for single-digit millisecond latencies. Observed latency is application specific.

NEW QUESTION 8

Your VMs are running in a subnet that has a subnet mask of 255.255.255.240. The current subnet has no more free IP addresses and you require an additional 10 IP addresses for new VMs. The existing and new VMs should all be able to reach each other without additional routes. What should you do?

- A. Use gcloud to expand the IP range of the current subnet.
- B. Delete the subnet, and recreate it using a wider range of IP addresses.
- C. Create a new project
- D. Use Shared VPC to share the current network with the new project.
- E. Create a new subnet with the same starting IP but a wider range to overwrite the current subnet.

Answer: A

Explanation:

<https://cloud.google.com/sdk/gcloud/reference/compute/networks/subnets/expand-ip-range>

gcloud compute networks subnets expand-ip-range - expand the IP range of a Compute Engine subnetwork gcloud compute networks subnets expand-ip-range NAME --prefix-length=PREFIX_LENGTH [--region=REGION] [GCLOUD_WIDE_FLAG ...]

NEW QUESTION 9

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 10

Your company is using Google Workspace to manage employee accounts. Anticipated growth will increase the number of personnel from 100 employees to 1,000

employees within 2 years. Most employees will need access to your company's Google Cloud account. The systems and processes will need to support 10x growth without performance degradation, unnecessary complexity, or security issues. What should you do?

- A. Migrate the users to Active Director
- B. Connect the Human Resources system to Active Director
- C. Turn on Google Cloud Directory Sync (GCDS) for Cloud Identity
- D. Turn on Identity Federation from Cloud Identity to Active Directory.
- E. Organize the users in Cloud Identity into group
- F. Enforce multi-factor authentication in Cloud Identity.
- G. Turn on identity federation between Cloud Identity and Google Workspac
- H. Enforce multi-factor authentication for domain wide delegation.
- I. Use a third-party identity provider service through federatio
- J. Synchronize the users from Google Workplace to the third-party provider in real time.

Answer: B

NEW QUESTION 10

You have a virtual machine that is currently configured with 2 vCPUs and 4 GB of memory. It is running out of memory. You want to upgrade the virtual machine to have 8 GB of memory. What should you do?

- A. Rely on live migration to move the workload to a machine with more memory.
- B. Use gcloud to add metadata to the V
- C. Set the key to required-memory-size and the value to 8 GB.
- D. Stop the VM, change the machine type to n1-standard-8, and start the VM.
- E. Stop the VM, increase the memory to 8 GB, and start the VM.

Answer: D

Explanation:

In Google compute engine, if predefined machine types don't meet your needs, you can create an instance with custom virtualized hardware settings. Specifically, you can create an instance with a custom number of vCPUs and custom memory, effectively using a custom machine type. Custom machine types are ideal for the following scenarios: 1. Workloads that aren't a good fit for the predefined machine types that are available you. 2. Workloads that require more processing power or more memory but don't need all of the upgrades that are provided by the next machine type level. In our scenario, we only need a memory upgrade. Moving to a bigger instance would also bump up the CPU which we don't need so we have to use a custom machine type. It is not possible to change memory while the instance is running so you need to first stop the instance, change the memory and then start it again. See below a screenshot that shows how CPU/Memory can be customized for an instance that has been stopped. Ref: <https://cloud.google.com/compute/docs/instances/creating-instance-with-custom-machine-type>

NEW QUESTION 12

Your web application has been running successfully on Cloud Run for Anthos. You want to evaluate an updated version of the application with a specific percentage of your production users (canary deployment). What should you do?

- A. Create a new service with the new version of the applicatio
- B. Split traffic between this version and the version that is currently running.
- C. Create a new revision with the new version of the applicatio
- D. Split traffic between this version and the version that is currently running.
- E. Create a new service with the new version of the applicatio
- F. Add an HTTP Load Balancer in front of both services.
- G. Create a new revision with the new version of the applicatio
- H. Add an HTTP Load Balancer in front of both revisions.

Answer: B

Explanation:

<https://cloud.google.com/kuberun/docs/rollouts-rollbacks-traffic-migration>

NEW QUESTION 15

You have a workload running on Compute Engine that is critical to your business. You want to ensure that the data on the boot disk of this workload is backed up regularly. You need to be able to restore a backup as quickly as possible in case of disaster. You also want older backups to be cleaned automatically to save on cost. You want to follow Google-recommended practices. What should you do?

- A. Create a Cloud Function to create an instance template.
- B. Create a snapshot schedule for the disk using the desired interval.
- C. Create a cron job to create a new disk from the disk using gcloud.
- D. Create a Cloud Task to create an image and export it to Cloud Storage.

Answer: B

Explanation:

Best practices for persistent disk snapshots

You can create persistent disk snapshots at any time, but you can create snapshots more quickly and with greater reliability if you use the following best practices.

Creating frequent snapshots efficiently

Use snapshots to manage your data efficiently.

Create a snapshot of your data on a regular schedule to minimize data loss due to unexpected failure. Improve performance by eliminating excessive snapshot downloads and by creating an image and reusing it. Set your snapshot schedule to off-peak hours to reduce snapshot time.

Snapshot frequency limits

Creating snapshots from persistent disks

You can snapshot your disks at most once every 10 minutes. If you want to issue a burst of requests to snapshot your disks, you can issue at most 6 requests in 60 minutes.

If the limit is exceeded, the operation fails and returns the following error: <https://cloud.google.com/compute/docs/disks/snapshot-best-practices>

NEW QUESTION 18

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 19

A colleague handed over a Google Cloud Platform project for you to maintain. As part of a security checkup, you want to review who has been granted the Project Owner role. What should you do?

- A. In the console, validate which SSH keys have been stored as project-wide keys.
- B. Navigate to Identity-Aware Proxy and check the permissions for these resources.
- C. Enable Audit Logs on the IAM & admin page for all resources, and validate the results.
- D. Use the command `gcloud projects get-iam-policy` to view the current role assignments.

Answer: D

Explanation:

A simple approach would be to use the command flags available when listing all the IAM policy for a given project. For instance, the following command: `gcloud projects get-iam-policy $PROJECT_ID`

```
--flatten="bindings[].members" --format="table(bindings.members)" --filter="bindings.role:roles/owner"
```

outputs all the users and service accounts associated with the role 'roles/owner' in the project in question. <https://groups.google.com/g/google-cloud-dev/c/Z6sZs7TvygQ?pli=1>

NEW QUESTION 21

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 to your instance.
- B. Create an alert in Cloud Monitoring to alert when the percentage of 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 identify 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 25

You are running multiple VPC-native Google Kubernetes Engine clusters in the same subnet. The IPs available for the nodes are exhausted, and you want to ensure that the clusters can grow in nodes when needed. What should you do?

- A. Create a new subnet in the same region as the subnet being used.
- B. Add an alias IP range to the subnet used by the GKE clusters.
- C. Create a new VPC, and set up VPC peering with the existing VPC.
- D. Expand the CIDR range of the relevant subnet for the cluster.

Answer: D

Explanation:

`gcloud compute networks subnets expand-ip-range NAME gcloud compute networks subnets expand-ip-range`

- expand the IP range of a Compute Engine subnetwork <https://cloud.google.com/sdk/gcloud/reference/compute/networks/subnets/expand-ip-range>

NEW QUESTION 29

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 30

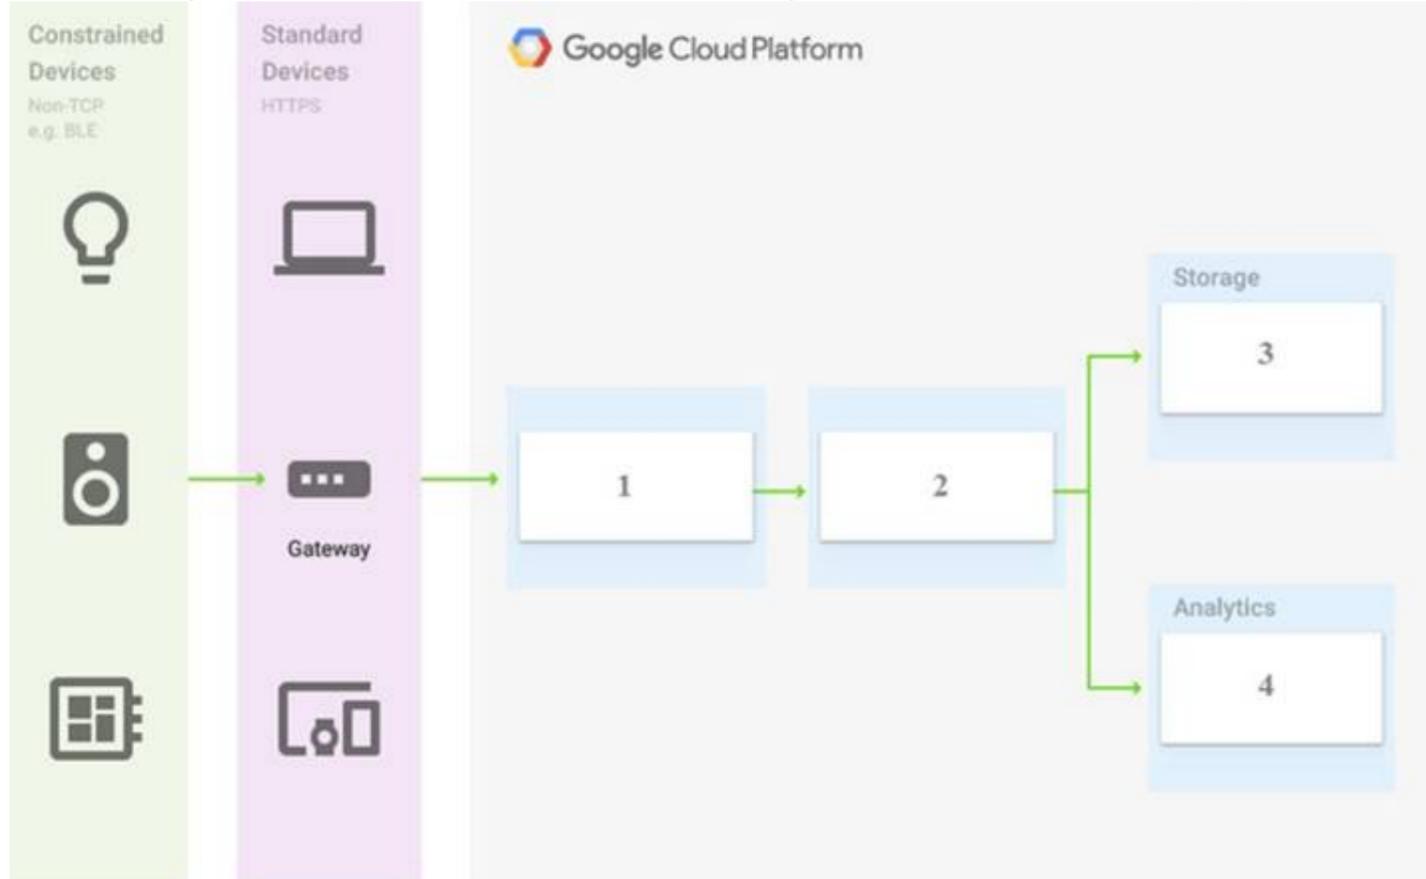
Your company uses Cloud Storage to store application backup files for disaster recovery purposes. You want to follow Google's recommended practices. Which storage option should you use?

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

Answer: D

NEW QUESTION 34

You are building a pipeline to process time-series data. Which Google Cloud Platform services should you put in boxes 1,2,3, and 4?



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

Answer: D

NEW QUESTION 37

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 40

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 45

Your application development team has created Docker images for an application that will be deployed on Google Cloud. Your team does not want to manage the infrastructure associated with this application. You need to ensure that the application can scale automatically as it gains popularity. What should you do?

- A. Create an Instance template with the container image, and deploy a Managed Instance Group with Autoscaling.
- B. Upload Docker images to Artifact Registry, and deploy the application on Google Kubernetes Engine using Standard mode.
- C. Upload Docker images to the Cloud Storage, and deploy the application on Google Kubernetes Engine using Standard mode.
- D. Upload Docker images to Artifact Registry, and deploy the application on Cloud Run.

Answer: D

NEW QUESTION 46

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

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

Answer: A

Explanation:

https://cloud.google.com/iam/docs/job-functions/auditing#scenario_external_auditors

Because if you directly add users to the IAM roles, then if any users left the organization then you have to remove the users from multiple places and need to revoke his/her access from multiple places. But, if you put a user into a group then its very easy to manage these type of situations. Now, if any user left then you just need to remove the user from the group and all the access got revoked

The organization creates a Google group for these external auditors and adds the current auditor to the group. This group is monitored and is typically granted access to the dashboard application. During normal access, the auditors' Google group is only granted access to view the historic logs stored in BigQuery. If any anomalies are discovered, the group is granted permission to view the actual Cloud Logging Admin Activity logs via the dashboard's elevated access mode. At the end of each audit period, the group's access is then revoked. Data is redacted using Cloud DLP before being made accessible for viewing via the dashboard application. The table below explains IAM logging roles that an Organization Administrator can grant to the service account used by the dashboard, as well as the resource level at which the role is granted.

NEW QUESTION 47

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

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

Answer: A

Explanation:

When setting support levels for permissions in custom roles, you can set to one of SUPPORTED, TESTING or NOT_SUPPORTED.

Ref: <https://cloud.google.com/iam/docs/custom-roles-permissions-support>

NEW QUESTION 50

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 55

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

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

Answer: C

NEW QUESTION 58

You need to enable traffic between multiple groups of Compute Engine instances that are currently running two different GCP projects. Each group of Compute Engine instances is running in its own VPC. What should you do?

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

Answer: B

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 62

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

- A. Create two configurations using gcloud confi
- B. Write a script that sets configurations as active, individuall
- C. For each configuration, use gcloud compute instances list to get a list of compute resources.
- D. Create two configurations using gsutil confi
- E. Write a script that sets configurations as active, individuall

- F. For each configuration, use `gsutil compute instances list` to get a list of compute resources.
- G. Go to Cloud Shell and export this information to Cloud Storage on a daily basis.
- H. Go to GCP Console and export this information to Cloud SQL on a daily basis.

Answer: A

Explanation:

You can create two configurations – one for the development project and another for the production project. And you do that by running “`gcloud config configurations create`” command. <https://cloud.google.com/sdk/gcloud/reference/config/configurations/create>In your custom script, you can load these configurations one at a time and execute `gcloud compute instances list` to list Google Compute Engine instances in the project that is active in the `gcloud` configuration. Ref: <https://cloud.google.com/sdk/gcloud/reference/compute/instances/list> Once you have this information, you can export it in a suitable format to a suitable target e.g. export as CSV or export to Cloud Storage/BigQuery/SQL, etc

NEW QUESTION 65

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 68

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

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

Answer: B

Explanation:

Creating scheduled snapshots for persistent disk This document describes how to create a snapshot schedule to regularly and automatically back up your zonal and regional persistent disks. Use snapshot schedules as a best practice to back up your Compute Engine workloads. After creating a snapshot schedule, you can apply it to one or more persistent disks. <https://cloud.google.com/compute/docs/disks/scheduled-snapshots>

NEW QUESTION 70

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 71

You need to configure optimal data storage for files stored in Cloud Storage for minimal cost. The files are used in a mission-critical analytics pipeline that is used continually. The users are in Boston, MA (United States). What should you do?

- A. Configure regional storage for the region closest to the users Configure a Nearline storage class
- B. Configure regional storage for the region closest to the users Configure a Standard storage class
- C. Configure dual-regional storage for the dual region closest to the users Configure a Nearline storage class
- D. Configure dual-regional storage for the dual region closest to the users Configure a Standard storage class

Answer: B

Explanation:

Keywords: - continually -> Standard - mission-critical analytics -> dual-regional

NEW QUESTION 75

You are building a new version of an application hosted in an App Engine environment. You want to test the new version with 1% of users before you completely switch your application over to the new version. What should you do?

- A. Deploy a new version of your application in Google Kubernetes Engine instead of App Engine and then use GCP Console to split traffic.
- B. Deploy a new version of your application in a Compute Engine instance instead of App Engine and then use GCP Console to split traffic.
- C. Deploy a new version as a separate app in App Engine
- D. Then configure App Engine using GCP Console to split traffic between the two apps.
- E. Deploy a new version of your application in App Engine
- F. Then go to App Engine settings in GCP Console and split traffic between the current version and newly deployed versions accordingly.

Answer: D

Explanation:

GCP App Engine natively offers traffic splitting functionality between versions. You can use traffic splitting to specify a percentage distribution of traffic across two or more of the versions within a service. Splitting traffic allows you to conduct A/B testing between your versions and provides control over the pace when rolling out features.

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

NEW QUESTION 80

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 82

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 VM
- D. Start and stop the instance as needed.
- E. Create an Instance Template with Preemptible VMs
- F. Create a Managed Instance Group from the template and adjust Target CPU Utilization
- 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 84

You need to deploy an application, which is packaged in a container image, in a new project. The application exposes an HTTP endpoint and receives very few requests per day. You want to minimize costs. What should you do?

- A. Deploy the container on Cloud Run.
- B. Deploy the container on Cloud Run on GKE.
- C. Deploy the container on App Engine Flexible.
- D. Deploy the container on Google Kubernetes Engine, with cluster autoscaling and horizontal pod autoscaling enabled.

Answer: A

Explanation:

Cloud Run takes any container images and pairs great with the container ecosystem: Cloud Build, Artifact Registry, Docker. ... No infrastructure to manage: once deployed, Cloud Run manages your services so you can sleep well. Fast autoscaling. Cloud Run automatically scales up or down from zero to N depending on traffic.

<https://cloud.google.com/run>

NEW QUESTION 85

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 INTO dataset.platform-logs (timestamp, log) SELECT timestamp, log FROM compute.logs WHERE 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 86

Your existing application running in Google Kubernetes Engine (GKE) consists of multiple pods running on four GKE n1-standard-2 nodes. You need to deploy additional pods requiring n2-highmem-16 nodes without any downtime. What should you do?

- A. Use `gcloud container clusters upgrad`
- B. Deploy the new services.
- C. Create a new Node Pool and specify machine type n2-highmem-16. Deploy the new pods.
- D. Create a new cluster with n2-highmem-16 node
- E. Redeploy the pods and delete the old cluster.
- F. Create a new cluster with both n1-standard-2 and n2-highmem-16 node
- G. Redeploy the pods and delete the old cluster.

Answer: B

Explanation:

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

NEW QUESTION 89

You significantly changed a complex Deployment Manager template and want to confirm that the dependencies of all defined resources are properly met before committing it to the project. You want the most rapid feedback on your changes. What should you do?

- A. Use granular logging statements within a Deployment Manager template authored in Python.
- B. Monitor activity of the Deployment Manager execution on the Stackdriver Logging page of the GCP Console.
- C. Execute the Deployment Manager template against a separate project with the same configuration, and monitor for failures.
- D. Execute the Deployment Manager template using the `--preview` option in the same project, and observe the state of interdependent resources.

Answer: D

NEW QUESTION 92

Your company has an internal application for managing transactional orders. The application is used exclusively by employees in a single physical location. The application requires strong consistency, fast queries, and ACID guarantees for multi-table transactional updates. The first version of the application is implemented in PostgreSQL, and you want to deploy it to the cloud with minimal code changes. Which database is most appropriate for this application?

- A. BigQuery
- B. Cloud SQL
- C. Cloud Spanner
- D. Cloud Datastore

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/postgres>

NEW QUESTION 93

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

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

Answer: A

Explanation:

Implied rules Every VPC network has two implied firewall rules. These rules exist, but are not shown in the Cloud Console: Implied allow egress rule. An egress rule whose action is allow, destination is 0.0.0.0/0, and priority is the lowest possible (65535) lets any instance send traffic to any destination, except for traffic blocked by Google Cloud. A higher priority firewall rule may restrict outbound access. Internet access is allowed if no other firewall rules deny outbound traffic and if the instance has an external IP address or uses a Cloud NAT instance. For more information, see Internet access requirements. Implied deny ingress rule. An ingress rule whose action is deny, source is 0.0.0.0/0, and priority is the lowest possible (65535) protects all instances by blocking incoming connections to them. A higher priority rule might allow incoming access. The default network includes some additional rules that override this one, allowing certain types of incoming connections. https://cloud.google.com/vpc/docs/firewalls#default_firewall_rules

NEW QUESTION 98

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 100

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 datase
- C. Configure the table expiration to 60 days.
- D. Create a Stackdriver Logging Export with a Sink destination to Cloud Storang
- 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 BigQuer
- 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 103

You created a cluster.YAML file containing

- resources:
- name: cluster
- type: container.v1.cluster
- properties:
- zone: europe-west1-b
- cluster:
- description: My GCP ACE cluster
- initialNodeCount: 2

You want to use Cloud Deployment Manager to create this cluster in GKE. What should you do?

- A. gcloud deployment-manager deployments create my-gcp-ace-cluster --config cluster.yaml
- B. gcloud deployment-manager deployments create my-gcp-ace-cluster --type container.v1.cluster --config cluster.yaml
- C. gcloud deployment-manager deployments apply my-gcp-ace-cluster --type container.v1.cluster --config cluster.yaml
- D. gcloud deployment-manager deployments apply my-gcp-ace-cluster --config cluster.yaml

Answer: D

Explanation:

gcloud deployment-manager deployments create creates deployments based on the configuration file. (Infrastructure as code). All the configuration related to the artifacts is in the configuration file. This command correctly creates a cluster based on the provided cluster.yaml configuration file.

Ref: <https://cloud.google.com/sdk/gcloud/reference/deployment-manager/deployments/create>

NEW QUESTION 108

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 111

You are deploying a production application on Compute Engine. You want to prevent anyone from accidentally destroying the instance by clicking the wrong button. What should you do?

- A. Disable the flag "Delete boot disk when instance is deleted."
- B. Enable delete protection on the instance.
- C. Disable Automatic restart on the instance.
- D. Enable Preemptibility on the instance.

Answer: D

Explanation:

Preventing Accidental VM Deletion This document describes how to protect specific VM instances from deletion by setting the deletionProtection property on an Instance resource. To learn more about VM instances, read the Instances documentation. 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. <https://cloud.google.com/compute/docs/instances/preventing-accidental-vm-deletion>

NEW QUESTION 115

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 desktop.
- B. Verify that a firewall rule for port 3389 exists.
- C. Install a RDP client in your desktop.
- D. Set a Windows username and password in the GCP Console.
- E. Use the credentials to log in to the instance.
- F. Set a Windows password in the GCP Console.
- G. Verify that a firewall rule for port 22 exists.
- 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 Console.
- J. Verify that a firewall rule for port 3389 exists.
- 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 120

You need to deploy an application in Google Cloud using serverless technology. You want to test a new version of the application with a small percentage of production traffic. What should you do?

- A. Deploy the application to Cloud Run.
- B. Use gradual rollouts for traffic splitting.
- C. Deploy the application to Google Kubernetes Engine.
- D. Use Anthos Service Mesh for traffic splitting.
- E. Deploy the application to Cloud Functions.
- F. Saucily the version number in the function name.
- G. Deploy the application to App Engine.

I. For each new version, create a new service.

Answer: A

NEW QUESTION 123

You have a Bigtable instance that consists of three nodes that store personally identifiable information (PII) data. You need to log all read or write operations, including any metadata or configuration reads of this database table, in your company's Security Information and Event Management (SIEM) system. What should you do?

- A. • Navigate to Cloud Monitoring in the Google Cloud console, and create a custom monitoring job for the Bigtable instance to track all changes. • Create an alert by using webhook endpoint
- B. with the SIEM endpoint as a receiver
- C. Navigate to the Audit Logs page in the Google Cloud console, and enable Data Read and Admin Read logs for the Bigtable instance
- D. Data Write and Admin Read logs for the Bigtable instance • Create a Pub/Sub topic as a Cloud Logging sink destination, and add your SIEM as a subscriber to the topic.
- E. • Install the Ops Agent on the Bigtable instance during configuration
- F. • Create a service account with read permissions for the Bigtable instance. • Create a custom Dataflow job with this service account to export logs to the company's SIEM system.
- G. • Navigate to the Audit Logs page in the Google Cloud console, and enable Admin Write logs for the Bigtable instance. • Create a Cloud Functions instance to export logs from Cloud Logging to your SIEM.

Answer: B

NEW QUESTION 124

You need to verify that a Google Cloud Platform service account was created at a particular time. What should you do?

- A. Filter the Activity log to view the Configuration category
- B. Filter the Resource type to Service Account.
- C. Filter the Activity log to view the Data Access category
- D. Filter the Resource type to Google Project.
- E. Filter the Activity log to view the Data Access category
- F. Filter the Resource type to Service Account.
- G. Filter the Activity log to view the Data Access category
- H. Filter the Resource type to Google Project.

Answer: A

Explanation:

<https://developers.google.com/cloud-search/docs/guides/audit-logging-manual>

NEW QUESTION 125

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

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

Answer: D

Explanation:

"The Cloud Spanner to Cloud Storage Text template is a batch pipeline that reads in data from a Cloud Spanner table, optionally transforms the data via a JavaScript User Defined Function (UDF) that you provide, and writes it to Cloud Storage as CSV text files."

<https://cloud.google.com/dataflow/docs/guides/templates/provided-batch#cloudspannertogcstext>

"The Dataflow connector for Cloud Spanner lets you read data from and write data to Cloud Spanner in a Dataflow pipeline"

<https://cloud.google.com/spanner/docs/dataflow-connector> <https://cloud.google.com/bigquery/external-data-sources>

NEW QUESTION 130

You are migrating a production-critical on-premises application that requires 96 vCPUs to perform its task. You want to make sure the application runs in a similar environment on GCP. What should you do?

- A. When creating the VM, use machine type n1-standard-96.
- B. When creating the VM, use Intel Skylake as the CPU platform.
- C. Create the VM using Compute Engine default settings
- D. Use gcloud to modify the running instance to have 96 vCPUs.
- E. Start the VM using Compute Engine default settings, and adjust as you go based on Rightsizing Recommendations.

Answer: A

Explanation:

Ref: https://cloud.google.com/compute/docs/machine-types#n1_machine_type

NEW QUESTION 135

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?

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

Answer: B

NEW QUESTION 138

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?

- Deployment Manager
- Cloud Composer
- Managed Instance Group
- Unmanaged Instance Group

Answer: A

Explanation:

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

NEW QUESTION 141

You have an instance group that you want to load balance. You want the load balancer to terminate the client SSL session. The instance group is used to serve a public web application over HTTPS. You want to follow Google-recommended practices. What should you do?

- Configure an HTTP(S) load balancer.
- Configure an internal TCP load balancer.
- Configure an external SSL proxy load balancer.
- Configure an external TCP proxy load balancer.

Answer: A

NEW QUESTION 145

You create a Deployment with 2 replicas in a Google Kubernetes Engine cluster that has a single preemptible node pool. After a few minutes, you use kubectl to examine the status of your Pod and observe that one of them is still in Pending status:

```
$ kubectl get pods -l app=myapp
NAME                                READY    STATUS    RESTART    AGE
myapp-deployment-58ddb995-1p86m     0/1     Pending  0          9m
myapp-deployment-58ddb995-qjpkg     1/1     Running  0          9m
```

What is the most likely cause?

- The pending Pod's resource requests are too large to fit on a single node of the cluster.
- Too many Pods are already running in the cluster, and there are not enough resources left to schedule the pending Pod.
- The node pool is configured with a service account that does not have permission to pull the container image used by the pending Pod.
- The pending Pod was originally scheduled on a node that has been preempted between the creation of the Deployment and your verification of the Pods' statu
- It is currently being rescheduled on a new node.

Answer: B

Explanation:

- The pending Pods resource requests are too large to fit on a single node of the cluster. Too many Pods are already running in the cluster, and there are not enough resources left to schedule the pending Pod. is the right answer.
- When you have a deployment with some pods in running and other pods in the pending state, more often than not it is a problem with resources on the nodes. Heres a sample output of this use case. We see that the problem is with insufficient CPU on the Kubernetes nodes so we have to either enable auto-scaling or manually scale up the nodes.

NEW QUESTION 147

You have an application that receives SSL-encrypted TCP traffic on port 443. Clients for this application are located all over the world. You want to minimize latency for the clients. Which load balancing option should you use?

- HTTPS Load Balancer
- Network Load Balancer
- SSL Proxy Load Balancer
- Internal TCP/UDP Load Balance
- Add a firewall rule allowing ingress traffic from 0.0.0.0/0 on the target instances.

Answer: C

NEW QUESTION 150

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

- A. Run a test using simulated maintenance event
- B. If the test is successful, use preemptible N1 Standard VMs when running future jobs.
- C. Run a test using simulated maintenance event
- D. If the test is successful, use N1 Standard VMs when running future jobs.
- E. Run a test using a managed instance group
- F. If the test is successful, use N1 Standard VMs in the managed instance group when running future jobs.
- G. Run a test using N1 standard VMs instead of N2. If the test is successful, use N1 Standard VMs when running future jobs.

Answer: A

Explanation:

Creating and starting a preemptible VM instance This page explains how to create and use a preemptible virtual machine (VM) instance. A preemptible instance is an instance you can create and run at a much lower price than normal instances. However, Compute Engine might terminate (preempt) these instances if it requires access to those resources for other tasks. Preemptible instances will always terminate after 24 hours. To learn more about preemptible instances, read the preemptible instances documentation. Preemptible instances are recommended only for fault-tolerant applications that can withstand instance preemptions. Make sure your application can handle preemptions before you decide to create a preemptible instance. To understand the risks and value of preemptible instances, read the preemptible instances documentation. <https://cloud.google.com/compute/docs/instances/create-start-preemptible-instance>

NEW QUESTION 153

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 156

The storage costs for your application logs have far exceeded the project budget. The logs are currently being retained indefinitely in the Cloud Storage bucket myapp-gcp-ace-logs. You have been asked to remove logs older than 90 days from your Cloud Storage bucket. You want to optimize ongoing Cloud Storage spend. What should you do?

- A. Write a script that runs `gsutil ls -l -gs://myapp-gcp-ace-logs/**` to find and remove items older than 90 days
- B. Schedule the script with cron.
- C. Write a lifecycle management rule in JSON and push it to the bucket with `gsutil lifecycle set config-json-file`.
- D. Write a lifecycle management rule in XML and push it to the bucket with `gsutil lifecycle set config-xml-file`.
- E. Write a script that runs `gsutil ls -lr gs://myapp-gcp-ace-logs/**` to find and remove items older than 90 days
- F. Repeat this process every morning.

Answer: B

Explanation:

You write a lifecycle management rule in XML and push it to the bucket with `gsutil lifecycle set config-xml-file`. is not right.

`gsutil lifecycle set` enables you to set the lifecycle configuration on one or more buckets based on the configuration file provided. However, XML is not a valid supported type for the configuration file.

Ref: <https://cloud.google.com/storage/docs/gsutil/commands/lifecycle>

➤ Write a script that runs `gsutil ls -lr gs://myapp-gcp-ace-logs/**` to find and remove items older than 90 days. Repeat this process every morning. is not right. This manual approach is error-prone, time-consuming and expensive. GCP Cloud Storage provides lifecycle management rules that let you achieve this with minimal effort.

➤ Write a script that runs `gsutil ls -l gs://myapp-gcp-ace-logs/**` to find and remove items older than 90 days. Schedule the script with cron. is not right. This manual approach is error-prone, time-consuming and expensive. GCP Cloud Storage provides lifecycle management rules that let you achieve this with minimal effort.

➤ Write a lifecycle management rule in JSON and push it to the bucket with `gsutil lifecycle set config-json-file`. is the right answer.

You can assign a lifecycle management configuration to a bucket. The configuration contains a set of rules which apply to current and future objects in the bucket. When an object meets the criteria of one of the rules, Cloud Storage automatically performs a specified action on the object. One of the supported actions is to Delete objects. You can set up a lifecycle management to delete objects older than 90 days. `gsutil lifecycle set` enables you to set the lifecycle configuration on the bucket based on the configuration file. JSON is the only supported type for the configuration file. The `config-json-file` specified on the command line should be a path to a local file containing the lifecycle configuration JSON document.

Ref: <https://cloud.google.com/storage/docs/gsutil/commands/lifecycle> Ref: <https://cloud.google.com/storage/docs/lifecycle>

NEW QUESTION 158

Your organization uses G Suite for communication and collaboration. All users in your organization have a G Suite account. You want to grant some G Suite users access to your Cloud Platform project. What should you do?

- A. Enable Cloud Identity in the GCP Console for your domain.
- B. Grant them the required IAM roles using their G Suite email address.
- C. Create a CSV sheet with all users' email addresses
- D. Use the gcloud command line tool to convert them into Google Cloud Platform accounts.
- E. In the G Suite console, add the users to a special group called `cloud-console-users@yourdomain.com`. Rely on the default behavior of the Cloud Platform to grant users access if they are members of this group.

Answer: B

NEW QUESTION 163

Your company requires all developers to have the same permissions, regardless of the Google Cloud project they are working on. Your company's security policy also restricts developer permissions to Compute Engine.

Cloud Functions, and Cloud SQL. You want to implement the security policy with minimal effort. What should you do?

- A. • Create a custom role with Compute Engine, Cloud Functions, and Cloud SQL permissions in one project within the Google Cloud organization. • Copy the role across all projects created within the organization with the `gcloud iam roles copy` command. • Assign the role to developers in those projects.
- B. • Add all developers to a Google group in Google Groups for Workspace. • Assign the predefined role of Compute Admin to the Google group at the Google Cloud organization level.
- C. • Add all developers to a Google group in Cloud Identity. • Assign predefined roles for Compute Engine, Cloud Functions, and Cloud SQL permissions to the Google group for each project in the Google Cloud organization.
- D. • Add all developers to a Google group in Cloud Identity. • Create a custom role with Compute Engine, Cloud Functions, and Cloud SQL permissions at the Google Cloud organization level. • Assign the custom role to the Google group.

Answer: D

Explanation:

<https://www.cloudskillsboost.google/focuses/1035?parent=catalog#:~:text=custom%20role%20at%20the%20or>

NEW QUESTION 165

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 167

You are working with a Cloud SQL MySQL database at your company. You need to retain a month-end copy of the database for three years for audit purposes. What should you do?

- A. Save file automatic first-of-the- month backup for three years Store the backup file in an Archive class Cloud Storage bucket
- B. Convert the automatic first-of-the-month backup to an export file Write the export file to a Coldline class Cloud Storage bucket
- C. Set up an export job for the first of the month Write the export file to an Archive class Cloud Storage bucket
- D. Set up an on-demand backup for the first of the month Write the backup to an Archive class Cloud Storage bucket

Answer: C

Explanation:

https://cloud.google.com/sql/docs/mysql/backup-recovery/backups#can_i_export_a_backup https://cloud.google.com/sql/docs/mysql/import-export#automating_export_operations

NEW QUESTION 172

Your company implemented BigQuery as an enterprise data warehouse. Users from multiple business units run queries on this data warehouse. However, you notice that query costs for BigQuery are very high, and you need to control costs. Which two methods should you use? (Choose two.)

- A. Split the users from business units to multiple projects.
- B. Apply a user- or project-level custom query quota for BigQuery data warehouse.
- C. Create separate copies of your BigQuery data warehouse for each business unit.
- D. Split your BigQuery data warehouse into multiple data warehouses for each business unit.
- E. Change your BigQuery query model from on-demand to flat rat
- F. Apply the appropriate number of slots to each Project.

Answer: BE

Explanation:

<https://cloud.google.com/bigquery/docs/custom-quotas> https://cloud.google.com/bigquery/pricing#flat_rate_pricing

NEW QUESTION 174

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 project
- 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 178

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

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

Answer: A

Explanation:

When we create subnets in the same VPC with different CIDR ranges, they can communicate automatically within VPC. Resources within a VPC network can communicate with one another by using internal (private) IPv4 addresses, subject to applicable network firewall rules
Ref: <https://cloud.google.com/vpc/docs/vpc>

NEW QUESTION 179

You want to verify the IAM users and roles assigned within a GCP project named my-project. What should you do?

- A. Run `gcloud iam roles list`
- B. Review the output section.
- C. Run `gcloud iam service-accounts list`
- D. Review the output section.
- E. Navigate to the project and then to the IAM section in the GCP Console
- F. Review the members and roles.
- G. Navigate to the project and then to the Roles section in the GCP Console
- H. Review the roles and status.

Answer: C

Explanation:

Logged onto console and followed the steps and was able to see all the assigned users and roles.

NEW QUESTION 182

You have a project for your App Engine application that serves a development environment. The required testing has succeeded and you want to create a new project to serve as your production environment. What should you do?

- A. Use `gcloud` to create the new project, and then deploy your application to the new project.
- B. Use `gcloud` to create the new project and to copy the deployed application to the new project.
- C. Create a Deployment Manager configuration file that copies the current App Engine deployment into a new project.
- D. Deploy your application again using `gcloud` and specify the project parameter with the new project name to create the new project.

Answer: A

Explanation:

You can deploy to a different project by using `--project` flag.

By default, the service is deployed the current project configured via:

```
$ gcloud config set core/project PROJECT
```

To override this value for a single deployment, use the `--project` flag:

```
$ gcloud app deploy ~/my_app/app.yaml --project=PROJECT Ref: https://cloud.google.com/sdk/gcloud/reference/app/deploy
```

NEW QUESTION 186

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

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

Answer: D

NEW QUESTION 190

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 195

You recently deployed a new version of an application to App Engine and then discovered a bug in the release. You need to immediately revert to the prior version of the application. What should you do?

- A. Run `gcloud app restore`.
- B. On the App Engine page of the GCP Console, select the application that needs to be reverted and click Revert.
- C. On the App Engine Versions page of the GCP Console, route 100% of the traffic to the previous version.
- D. Deploy the original version as a separate application.
- E. Then go to App Engine settings and split traffic between applications so that the original version serves 100% of the requests.

Answer: C

NEW QUESTION 200

You have created a new project in Google Cloud through the `gcloud` command line interface (CLI) and linked a billing account. You need to create a new Compute Engine instance using the CLI. You need to perform the prerequisite steps. What should you do?

- A. Create a Cloud Monitoring Workspace.
- B. Create a VPC network in the project.
- C. Enable the `compute.googleapis.com` API.
- D. Grant yourself the IAM role of Compute Admin.

Answer: D

NEW QUESTION 201

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 202

You have a managed instance group comprised of preemptible VM's. All of the VM's keep deleting and recreating themselves every minute. What is a possible cause of this behavior?

- A. Your zonal capacity is limited, causing all preemptible VM's to be shutdown to recover capacity.
- B. Try deploying your group to another zone.
- C. You have hit your instance quota for the region.
- D. Your managed instance group's VM's are toggled to only last 1 minute in preemptible settings.
- E. Your managed instance group's health check is repeatedly failing, either to a misconfigured health check or misconfigured firewall rules not allowing the health check to access the instance.

Answer: D

Explanation:

as the instances (normal or preemptible) would be terminated and relaunched if the health check fails either due to application not configured properly or the instances firewall do not allow health check to happen.

GCP provides health check systems that connect to virtual machine (VM) instances on a configurable, periodic basis. Each connection attempt is called a probe.

GCP records the success or failure of each probe.

Health checks and load balancers work together. Based on a configurable number of sequential successful or failed probes, GCP computes an overall health state for each VM in the load balancer. VMs that respond successfully for the configured number of times are considered healthy. VMs that fail to respond successfully for a separate number of times are unhealthy.

GCP uses the overall health state of each VM to determine its eligibility for receiving new requests. In addition to being able to configure probe frequency and

health state thresholds, you can configure the criteria that define a successful probe.

NEW QUESTION 204

You want to set up a Google Kubernetes Engine cluster. Verifiable node identity and integrity are required for the cluster, and nodes cannot be accessed from the internet. You want to reduce the operational cost of managing your cluster, and you want to follow Google-recommended practices. What should you do?

- A. Deploy a private autopilot cluster
- B. Deploy a public autopilot cluster.
- C. Deploy a standard public cluster and enable shielded nodes.
- D. Deploy a standard private cluster and enable shielded nodes.

Answer: D

NEW QUESTION 208

Your company set up a complex organizational structure on Google Cloud Platform. The structure includes hundreds of folders and projects. Only a few team members should be able to view the hierarchical structure. You need to assign minimum permissions to these team members and you want to follow Google-recommended practices. What should you do?

- A. Add the users to roles/browser role.
- B. Add the users to roles/iam.roleViewer role.
- C. Add the users to a group, and add this group to roles/browser role.
- D. Add the users to a group, and add this group to roles/iam.roleViewer role.

Answer: C

Explanation:

We need to apply the GCP Best practices. roles/browser Browser Read access to browse the hierarchy for a project, including the folder, organization, and IAM policy. This role doesn't include permission to view resources in the project. <https://cloud.google.com/iam/docs/understanding-roles>

NEW QUESTION 212

.....

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](#)