

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer

<https://www.2passeasy.com/dumps/Associate-Cloud-Engineer/>



NEW QUESTION 1

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

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

Answer: B

Explanation:

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

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

Node auto-upgrades help you keep the nodes in your cluster up to date with the cluster master version when your master is updated on your behalf. When you create a new cluster or node pool with Google Cloud Console or the gcloud command, node auto-upgrade is enabled by default.

Ref: <https://cloud.google.com/kubernetes-engine/docs/how-to/node-auto-upgrades>

NEW QUESTION 2

You have developed an application that consists of multiple microservices, with each microservice packaged in its own Docker container image. You want to deploy the entire application on Google Kubernetes Engine so that each microservice can be scaled individually. What should you do?

- A. Create and deploy a Custom Resource Definition per microservice.
- B. Create and deploy a Docker Compose File.
- C. Create and deploy a Job per microservice.
- D. Create and deploy a Deployment per microservice.

Answer: D

NEW QUESTION 3

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 4

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 5

Your company has a single sign-on (SSO) identity provider that supports Security Assertion Markup Language (SAML) integration with service providers. Your company has users in Cloud Identity. You would like users to authenticate using your company's SSO provider. What should you do?

- A. In Cloud Identity, set up SSO with Google as an identity provider to access custom SAML apps.

- B. In Cloud Identity, set up SSO with a third-party identity provider with Google as a service provider.
- C. Obtain OAuth 2.0 credentials, configure the user consent screen, and set up OAuth 2.0 for Mobile & Desktop Apps.
- D. Obtain OAuth 2.0 credentials, configure the user consent screen, and set up OAuth 2.0 for Web Server Applications.

Answer: B

Explanation:

https://support.google.com/cloudidentity/answer/6262987?hl=en&ref_topic=7558767

NEW QUESTION 6

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

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

Answer: C

Explanation:

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

NEW QUESTION 7

You need to set up permissions for a set of Compute Engine instances to enable them to write data into a particular Cloud Storage bucket. You want to follow Google-recommended practices. What should you do?

- A. Create a service account with an access scop
- B. Use the access scope 'https://www.googleapis.com/auth/devstorage.write_only'.
- C. Create a service account with an access scop
- D. Use the access scope 'https://www.googleapis.com/auth/cloud-platform'.
- E. Create a service account and add it to the IAM role 'storage.objectCreator' for that bucket.
- F. Create a service account and add it to the IAM role 'storage.objectAdmin' for that bucket.

Answer: C

Explanation:

https://cloud.google.com/iam/docs/understanding-service-accounts#using_service_accounts_with_compute_eng <https://cloud.google.com/storage/docs/access-control/iam-roles>

NEW QUESTION 8

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?

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

Answer: C

NEW QUESTION 9

You host a static website on Cloud Storage. Recently, you began to include links to PDF files on this site. Currently, when users click on the links to these PDF files, their browsers prompt them to save the file onto their local system. Instead, you want the clicked PDF files to be displayed within the browser window directly, without prompting the user to save the file locally. What should you do?

- A. Enable Cloud CDN on the website frontend.
- B. Enable 'Share publicly' on the PDF file objects.
- C. Set Content-Type metadata to application/pdf on the PDF file objects.
- D. Add a label to the storage bucket with a key of Content-Type and value of application/pdf.

Answer: C

Explanation:

https://developer.mozilla.org/en-US/docs/Web/HTTP/Basics_of_HTTP/MIME_Types#importance_of_setting_t

NEW QUESTION 10

You are asked to set up application performance monitoring on Google Cloud projects A, B, and C as a single pane of glass. You want to monitor CPU, memory, and disk. What should you do?

- A. Enable API and then share charts from project A, B, and C.
- B. Enable API and then give the metrics.reader role to projects A, B, and C.
- C. Enable API and then use default dashboards to view all projects in sequence.
- D. Enable API, create a workspace under project A, and then add project B and C.

Answer: D

Explanation:

<https://cloud.google.com/monitoring/settings/multiple-projects> <https://cloud.google.com/monitoring/workspaces>

NEW QUESTION 10

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 11

Your company has a 3-tier solution running on Compute Engine. The configuration of the current infrastructure is shown below.



Each tier has a service account that is associated with all instances within it. You need to enable communication on TCP port 8080 between tiers as follows:

- Instances in tier #1 must communicate with tier #2.
- Instances in tier #2 must communicate with tier #3. What should you do?

- A. 1. Create an ingress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.2.0/24)• Protocols: allow all2. Create an ingress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.1.0/24)•Protocols: allow all
- B. 1. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #2 service account• Source filter: all instances with tier #1 service account• Protocols: allow TCP:80802. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #3 service account• Source filter: all instances with tier #2 service account• Protocols: allow TCP: 8080
- C. 1. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #2 service account• Source filter: all instances with tier #1 service account• Protocols: allow all2. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #3 service account• Source filter: all instances with tier #2 service account• Protocols: allow all
- D. 1. Create an egress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.2.0/24)• Protocols: allow TCP: 80802. Create an egress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.1.0/24)• Protocols: allow TCP: 8080

Answer: B

Explanation:

* 1. Create an ingress firewall rule with the following settings: "⌘ Targets: all instances with tier #2 service account "⌘ Source filter: all instances with tier #1 service

account "⌘ Protocols: allow TCP:8080 2. Create an ingress firewall rule with the following settings: "⌘ Targets: all instances with tier #3 service account "⌘ Source filter: all instances with tier #2 service account "⌘ Protocols: allow TCP: 8080

NEW QUESTION 16

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-scope>

NEW QUESTION 18

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] [GLOUD_WIDE_FLAG ...]

NEW QUESTION 23

You need to immediately change the storage class of an existing Google Cloud bucket. You need to reduce service cost for infrequently accessed files stored in that bucket and for all files that will be added to that bucket in the future. What should you do?

- A. Use the gsutil to rewrite the storage class for the bucket Change the default storage class for the bucket
- B. Use the gsutil to rewrite the storage class for the bucket Set up Object Lifecycle management on the bucket
- C. Create a new bucket and change the default storage class for the bucket Set up Object Lifecycle management on lite bucket
- D. Create a new bucket and change the default storage class for the bucket import the files from the previous bucket into the new bucket

Answer: B

NEW QUESTION 28

Your company is moving its entire workload to Compute Engine. Some servers should be accessible through the Internet, and other servers should only be accessible over the internal network. All servers need to be able to talk to each other over specific ports and protocols. The current on-premises network relies on a demilitarized zone (DMZ) for the public servers and a Local Area Network (LAN) for the private servers. You need to design the networking infrastructure on Google Cloud to match these requirements. What should you do?

- A. 1. Create a single VPC with a subnet for the DMZ and a subnet for the LA
- B. 2. Set up firewall rules to open up relevant traffic between the DMZ and the LAN subnets, and another firewall rule to allow public ingress traffic for the DMZ.
- C. 1. Create a single VPC with a subnet for the DMZ and a subnet for the LA
- D. 2. Set up firewall rules to open up relevant traffic between the DMZ and the LAN subnets, and another firewall rule to allow public egress traffic for the DMZ.
- E. 1. Create a VPC with a subnet for the DMZ and another VPC with a subnet for the LA
- F. 2. Set up firewall rules to open up relevant traffic between the DMZ and the LAN subnets, and another firewall rule to allow public ingress traffic for the DMZ.
- G. 1. Create a VPC with a subnet for the DMZ and another VPC with a subnet for the LA
- H. 2. Set up firewall rules to open up relevant traffic between the DMZ and the LAN subnets, and another firewall rule to allow public egress traffic for the DMZ.

Answer: A

Explanation:

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

NEW QUESTION 30

You are developing a new application and are looking for a Jenkins installation to build and deploy your source code. You want to automate the installation as quickly and easily as possible. What should you do?

- A. Deploy Jenkins through the Google Cloud Marketplace.
- B. Create a new Compute Engine instance.
- C. Run the Jenkins executable.
- D. Create a new Kubernetes Engine cluster.
- E. Create a deployment for the Jenkins image.
- F. Create an instance template with the Jenkins executable.
- G. Create a managed instance group with this template.

Answer: A

Explanation:

Installing Jenkins

In this section, you use Cloud Marketplace to provision a Jenkins instance. You customize this instance to use the agent image you created in the previous section.

Go to the Cloud Marketplace solution for Jenkins. Click Launch on Compute Engine.

Change the Machine Type field to 4 vCPUs 15 GB Memory, n1-standard-4. Machine type selection for Jenkins deployment.

Click Deploy and wait for your Jenkins instance to finish being provisioned. When it is finished, you will see: Jenkins has been deployed.

https://cloud.google.com/solutions/using-jenkins-for-distributed-builds-on-compute-engine#installing_jenkins

NEW QUESTION 33

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 35

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 40

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 41

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 43

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 48

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

- A. Assign appropriate access for Google services to the service account used by the Compute Engine VM.
- B. Create a service account with appropriate access for Google services, and configure the application to use this account.
- C. Store credentials for service accounts with appropriate access for Google services in a config file, and deploy this config file with your application.
- D. Store credentials for your user account with appropriate access for Google services in a config file, and deploy this config file with your application.

Answer: B

Explanation:

In general, Google recommends that each instance that needs to call a Google API should run as a service account with the minimum permissions necessary for that instance to do its job. In practice, this means you should configure service accounts for your instances with the following process: Create a new service account rather than using the Compute Engine default service account. Grant IAM roles to that service account for only the resources that it needs. Configure the instance to run as that service account. Grant the instance the <https://www.googleapis.com/auth/cloud-platform> scope to allow full access to all Google Cloud APIs, so that the IAM permissions of the instance are completely determined by the IAM roles of the service account. Avoid granting more access than necessary and regularly check your service account permissions to make sure they are up-to-date.

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

NEW QUESTION 50

You are building an archival solution for your data warehouse and have selected Cloud Storage to archive your data. Your users need to be able to access this archived data once a quarter for some regulatory requirements. You want to select a cost-efficient option. Which storage option should you use?

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

Answer: A

Explanation:

Coldline Storage is a very-low-cost, highly durable storage service for storing infrequently accessed data. Coldline Storage is ideal for data you plan to read or modify at most once a quarter. Since we have a requirement to access data once a quarter and want to go with the most cost-efficient option, we should select Coldline Storage.

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

Google Cloud Storage Classes in the Organization

This slide represents the different types of storage classes such as multi-regional, regional, storage nearline, and storage cold line of the Google Cloud.

Storage Class	Characteristics	Use Cases	Price (Per Gb Per Month)*
Multi-Regional Storage	<ul style="list-style-type: none"> 99.95% availability Geo-redundant 	Keeps information that is frequently accessed around the globe, such as videos, gaming, and mobile applications	\$0.026 per GB/Month
Regional Storage	<ul style="list-style-type: none"> 99.9% availability Low cost per GB stored Data storage in a small region 	Keeps information that is frequently accessed around the globe, such as videos, gaming, and mobile applications	\$0.02 per GB/Month
Storage Nearline	<ul style="list-style-type: none"> 99.0% availability Very low cost per GB Data fetching costs Higher per-task costs 30-day minimum storage duration 	Keeps data that is not accessed is often ideal for data backups	\$0.01 per GB/Month
Storage Cold line	<ul style="list-style-type: none"> 99.0% availability Lowest cost per GB Data fetching costs Higher per-task costs 90-day minimum storage duration 	Keeps information that is infrequently ideal for disaster recovery or archived data	\$0.007 per GB/Month

This slide is 100% editable. Adapt it to your needs and capture your audience's attention.

NEW QUESTION 55

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: C

Explanation:

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

NEW QUESTION 56

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 60

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 61

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?

- A. Configure an HTTP(S) load balancer.
- B. Configure an internal TCP load balancer.
- C. Configure an external SSL proxy load balancer.
- D. Configure an external TCP proxy load balancer.

Answer: A

NEW QUESTION 66

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 68

You are using Google Kubernetes Engine with autoscaling enabled to host a new application. You want to expose this new application to the public, using HTTPS on a public IP address. What should you do?

- A. Create a Kubernetes Service of type NodePort for your application, and a Kubernetes Ingress to expose this Service via a Cloud Load Balancer.
- B. Create a Kubernetes Service of type ClusterIP for your applicatio
- C. Configure the public DNS name of your application using the IP of this Service.
- D. Create a Kubernetes Service of type NodePort to expose the application on port 443 of each node of the Kubernetes cluste
- E. Configure the public DNS name of your application with the IP of every node of the cluster to achieve load-balancing.
- F. Create a HAProxy pod in the cluster to load-balance the traffic to all the pods of the application.Forward the public traffic to HAProxy with an iptable rul
- G. Configure the DNS name of your application using the public IP of the node HAProxy is running on.

Answer: A

NEW QUESTION 70

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 75

You need to produce a list of the enabled Google Cloud Platform APIs for a GCP project using the gcloud command line in the Cloud Shell. The project name is my-project. What should you do?

- A. Run `gcloud projects list` to get the project ID, and then run `gcloud services list --project <project ID>`.
- B. Run `gcloud init` to set the current project to my-project, and then run `gcloud services list --available`.
- C. Run `gcloud info` to view the account value, and then run `gcloud services list --account <Account>`.
- D. Run `gcloud projects describe <project ID>` to verify the project value, and then run `gcloud services list--available`.

Answer: A

Explanation:

`gcloud services list --available` returns not only the enabled services in the project but also services that CAN be enabled.`

`https://cloud.google.com/sdk/gcloud/reference/services/list#--available`

Run the following command to list the enabled APIs and services in your current project: `gcloud services list`

whereas, Run the following command to list the APIs and services available to you in your current project: `gcloud services list --available`

`https://cloud.google.com/sdk/gcloud/reference/services/list#--available`

`--available`

Return the services available to the project to enable. This list will include any services that the project has already enabled.

To list the services the current project has enabled for consumption, run: `gcloud services list --enabled`

To list the services the current project can enable for consumption, run: `gcloud services list --available`

NEW QUESTION 79

You have created an application that is packaged into a Docker image. You want to deploy the Docker image as a workload on Google Kubernetes Engine. What should you do?

- A. Upload the image to Cloud Storage and create a Kubernetes Service referencing the image.
- B. Upload the image to Cloud Storage and create a Kubernetes Deployment referencing the image.
- C. Upload the image to Container Registry and create a Kubernetes Service referencing the image.
- D. Upload the image to Container Registry and create a Kubernetes Deployment referencing the image.

Answer: D

Explanation:

A deployment is responsible for keeping a set of pods running. A service is responsible for enabling network access to a set of pods.

NEW QUESTION 83

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 86

Your team maintains the infrastructure for your organization. The current infrastructure requires changes. You need to share your proposed changes with the rest of the team. You want to follow Google's recommended best practices. What should you do?

- A. Use Deployment Manager templates to describe the proposed changes and store them in a Cloud Storage bucket.
- B. Use Deployment Manager templates to describe the proposed changes and store them in Cloud Source Repositories.
- C. Apply the change in a development environment, run `gcloud compute instances list`, and then save the output in a shared Storage bucket.
- D. Apply the change in a development environment, run `gcloud compute instances list`, and then save the output in Cloud Source Repositories.

Answer: B

Explanation:

Showing Deployment Manager templates to your team will allow you to define the changes you want to implement in your cloud infrastructure. You can use Cloud Source Repositories to store Deployment Manager templates and collaborate with your team. Cloud Source Repositories are fully-featured, scalable, and private Git repositories you can use to store, manage and track changes to your code.

<https://cloud.google.com/source-repositories/docs/features>

NEW QUESTION 87

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

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

Answer: A

Explanation:

IAP controls access to your App Engine apps and Compute Engine VMs running on Google Cloud. It leverages user identity and the context of a request to determine if a user should be allowed access. IAP is a building block toward BeyondCorp, an enterprise security model that enables employees to work from untrusted networks without using a VPN.

By default, IAP uses Google identities and IAM. By leveraging Identity Platform instead, you can authenticate users with a wide range of external identity providers, such as:

Email/password

OAuth (Google, Facebook, Twitter, GitHub, Microsoft, etc.) SAML

OIDC

Phone number Custom Anonymous

This is useful if your application is already using an external authentication system, and migrating your users to Google accounts is impractical.

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

NEW QUESTION 88

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 89

You are deploying an application to a Compute Engine VM in a managed instance group. The application must be running at all times, but only a single instance of the VM should run per GCP project. How should you configure the instance group?

- A. Set autoscaling to On, set the minimum number of instances to 1, and then set the maximum number of instances to 1.
- B. Set autoscaling to Off, set the minimum number of instances to 1, and then set the maximum number of instances to 1.
- C. Set autoscaling to On, set the minimum number of instances to 1, and then set the maximum number of instances to 2.
- D. Set autoscaling to Off, set the minimum number of instances to 1, and then set the maximum number of instances to 2.

Answer: A

Explanation:

<https://cloud.google.com/compute/docs/autoscaler#specifications>

Autoscaling works independently from autohealing. If you configure autohealing for your group and an instance fails the health check, the autohealer attempts to recreate the instance. Recreating an instance can cause the number of instances in the group to fall below the autoscaling threshold (minNumReplicas) that you specify.

- > Since we need the application running at all times, we need a minimum 1 instance.
- > Only a single instance of the VM should run, we need a maximum 1 instance.
- > We want the application running at all times. If the VM crashes due to any underlying hardware failure, we want another instance to be added to MIG so that application can continue to serve requests. We can achieve this by enabling autoscaling. The only option that satisfies these three is Set autoscaling to On, set the minimum number of instances to 1, and then set the maximum number of instances to 1.

Ref: <https://cloud.google.com/compute/docs/autoscaler>

NEW QUESTION 90

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

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

Answer: A

Explanation:

https://cloud.google.com/resource-manager/docs/project-migration#oauth_consent_screen <https://cloud.google.com/resource-manager/docs/project-migration>

NEW QUESTION 92

You are about to deploy a new Enterprise Resource Planning (ERP) system on Google Cloud. The application holds the full database in-memory for fast data access, and you need to configure the most appropriate resources on Google Cloud for this application. What should you do?

- A. Provision preemptible Compute Engine instances.
- B. Provision Compute Engine instances with GPUs attached.
- C. Provision Compute Engine instances with local SSDs attached.
- D. Provision Compute Engine instances with M1 machine type.

Answer: D

Explanation:

M1 machine series Medium in-memory databases such as SAP HANA Tasks that require intensive use of memory with higher memory-to-vCPU ratios than the general-purpose high-memory machine types.

In-memory databases and in-memory analytics, business warehousing (BW) workloads, genomics analysis, SQL analysis services. Microsoft SQL Server and similar databases.

<https://cloud.google.com/compute/docs/machine-types>

<https://cloud.google.com/compute/docs/machine-types#:~:text=databases%20such%20as-,SAP%20HANA,-In%20memory,-database%3F>

<https://www.sap.com/india/products/hana.html#:~:text=is%20SAP%20HANA-,in%20memory,-database%3F>

NEW QUESTION 96


A company wants to build an application that stores images in a Cloud Storage bucket and wants to generate thumbnails as well as resize the images. They want to use a google managed service that can scale up and scale down to zero automatically with minimal effort. You have been asked to recommend a service. Which GCP service would you suggest?

- A. Google Compute Engine
- B. Google App Engine
- C. Cloud Functions
- D. Google Kubernetes Engine

Answer: C

Explanation:

Text Description automatically generated with low confidence



Cloud Functions is Google Cloud's event-driven serverless compute platform. It automatically scales based on the load and requires no additional configuration. You pay only for the resources used.

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

While all other options i.e. Google Compute Engine, Google Kubernetes Engine, Google App Engine support autoscaling, it needs to be configured explicitly based

on the load and is not as trivial as the scale up or scale down offered by Google's cloud functions.

NEW QUESTION 101

You want to select and configure a solution for storing and archiving data on Google Cloud Platform. You need to support compliance objectives for data from one geographic location. This data is archived after 30 days and needs to be accessed annually. What should you do?

- A. Select Multi-Regional Storage
- B. Add a bucket lifecycle rule that archives data after 30 days to Coldline Storage.
- C. Select Multi-Regional Storage
- D. Add a bucket lifecycle rule that archives data after 30 days to Nearline Storage.
- E. Select Regional Storage
- F. Add a bucket lifecycle rule that archives data after 30 days to Nearline Storage.
- G. Select Regional Storage
- H. Add a bucket lifecycle rule that archives data after 30 days to Coldline Storage.

Answer: D

Explanation:

Google Cloud Coldline is a new cold-tier storage for archival data with access frequency of less than once per year. Unlike other cold storage options, Nearline has no delays prior to data access, so now it is the leading solution among competitors.

The Real description is about Coldline storage Class: Coldline Storage

Coldline Storage is a very-low-cost, highly durable storage service for storing infrequently accessed data. Coldline Storage is a better choice than Standard Storage or Nearline Storage in scenarios where slightly lower availability, a 90-day minimum storage duration, and higher costs for data access are acceptable trade-offs for lowered at-rest storage costs.

Coldline Storage is ideal for data you plan to read or modify at most once a quarter. Note, however, that for data being kept entirely for backup or archiving purposes, Archive Storage is more cost-effective, as it offers the lowest storage costs.

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

NEW QUESTION 105

Your auditor wants to view your organization's use of data in Google Cloud. The auditor is most interested in auditing who accessed data in Cloud Storage buckets. You need to help the auditor access the data they need. What should you do?

- A. Assign the appropriate permissions, and then use Cloud Monitoring to review metrics
- B. Use the export logs API to provide the Admin Activity Audit Logs in the format they want
- C. Turn on Data Access Logs for the buckets they want to audit, and Then build a query in the log viewer that filters on Cloud Storage
- D. Assign the appropriate permissions, and then create a Data Studio report on Admin Activity Audit Logs

Answer: C

Explanation:

Types of audit logs Cloud Audit Logs provides the following audit logs for each Cloud project, folder, and organization: Admin Activity audit logs Data Access audit logs System Event audit logs Policy Denied audit logs ***Data Access audit logs contain API calls that read the configuration or metadata of resources, as well as user-driven API calls that create, modify, or read user-provided resource data. <https://cloud.google.com/logging/docs/audit#types>

<https://cloud.google.com/logging/docs/audit#data-access> Cloud Storage: When Cloud Storage usage logs are enabled, Cloud Storage writes usage data to the Cloud Storage bucket, which generates Data Access audit logs for the bucket. The generated Data Access audit log has its caller identity redacted.

NEW QUESTION 109

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 114

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 projec
- 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 117

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

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

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

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

Answer: B

Explanation:

<https://cloud.google.com/config-connector/docs/how-to/secrets#gcloud>

NEW QUESTION 120

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 125

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 127

Your company has a large quantity of unstructured data in different file formats. You want to perform ETL transformations on the data. You need to make the data accessible on Google Cloud so it can be processed by a Dataflow job. What should you do?

- A. Upload the data to BigQuery using the bq command line tool.
- B. Upload the data to Cloud Storage using the gsutil command line tool.
- C. Upload the data into Cloud SQL using the import function in the console.
- D. Upload the data into Cloud Spanner using the import function in the console.

Answer: B

Explanation:

"large quantity" : Cloud Storage or BigQuery "files" a file is nothing but an Object

NEW QUESTION 131

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 133

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 138

Your company uses a large number of Google Cloud services centralized in a single project. All teams have specific projects for testing and development. The DevOps team needs access to all of the production services in order to perform their job. You want to prevent Google Cloud product changes from broadening their permissions in the future. You want to follow Google-recommended practices. What should you do?

- A. Grant all members of the DevOps team the role of Project Editor on the organization level.
- B. Grant all members of the DevOps team the role of Project Editor on the production project.
- C. Create a custom role that combines the required permission
- D. Grant the DevOps team the custom role on the production project.
- E. Create a custom role that combines the required permission
- F. Grant the DevOps team the custom role on the organization level.

Answer: C

Explanation:

Understanding IAM custom roles

Key Point: 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.

Basic concepts

Custom roles are user-defined, and allow you to bundle one or more supported permissions to meet your specific needs. Custom roles are not maintained by Google; when new permissions, features, or services are added to Google Cloud, your custom roles will not be updated automatically.

When you create a custom role, you must choose an organization or project to create it in. You can then grant the custom role on the organization or project, as well as any resources within that organization or project.

https://cloud.google.com/iam/docs/understanding-custom-roles#basic_concepts

NEW QUESTION 142

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

- A. Deploy the container on Cloud Run for Anthos, and set the minimum number of instances to zero
- B. Deploy the container on Cloud Run (fully managed), and set the minimum number of instances to zero.
- C. Deploy the container on App Engine flexible environment with autoscalin

- D. and set the value min_instances to zero in the app yaml
- E. Deploy the container on App Engine flexible environment with manual scaling, and set the value instances to zero in the app yaml

Answer: B

Explanation:

https://cloud.google.com/kubernetes/docs/architecture-overview#components_in_the_default_installation

NEW QUESTION 147

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 billing 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 148

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

- A. Create an HTTP load balancer with a backend configuration that references an existing instance group. Set the health check to healthy (HTTP).
- B. Create an HTTP load balancer with a backend configuration that references an existing instance group. Define a balancing mode and set the maximum RPS to 10.
- C. Create a managed instance group
- D. Set the Autohealing health check to healthy (HTTP).
- E. Create a managed instance group
- 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 isn't 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 151

You are assigned to maintain a Google Kubernetes Engine (GKE) cluster named dev that was deployed on Google Cloud. You want to manage the GKE configuration using the command line interface (CLI). You have just downloaded and installed the Cloud SDK. You want to ensure that future CLI commands by default address this specific cluster. What should you do?

- A. Use the command `gcloud config set container/cluster dev`.
- B. Use the command `gcloud container clusters update dev`.
- C. Create a file called `gke.default` in the `~/.gcloud` folder that contains the cluster name.
- D. Create a file called `defaults.json` in the `~/.gcloud` folder that contains the cluster name.

Answer: A

Explanation:

To set a default cluster for gcloud commands, run the following command: `gcloud config set container/cluster CLUSTER_NAME`

<https://cloud.google.com/kubernetes-engine/docs/how-to/managing-clusters?hl=en>

NEW QUESTION 156

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 161

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 164

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 projec
- B. Update the existing project to link it to the existing billing account.
- C. Verify that you are Project Billing Manager for the GCP projec
- 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 accoun
- 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 accoun
- 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 166

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

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

Answer: A

Explanation:

Keyword, Version, traffic splitting, App Engine supports traffic splitting for versions before releasing.

NEW QUESTION 169

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

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

Answer: 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 170

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 175

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 176

You need to run an important query in BigQuery but expect it to return a lot of records. You want to find out how much it will cost to run the query. You are using on-demand pricing. What should you do?

- A. Arrange to switch to Flat-Rate pricing for this query, then move back to on-demand.
- B. Use the command line to run a dry run query to estimate the number of bytes read
- C. Then convert that bytes estimate to dollars using the Pricing Calculator.
- D. Use the command line to run a dry run query to estimate the number of bytes returned
- E. Then convert that bytes estimate to dollars using the Pricing Calculator.
- F. Run a select count (*) to get an idea of how many records your query will look through
- G. Then convert that number of rows to dollars using the Pricing Calculator.

Answer: B

NEW QUESTION 179

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 Configuration 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 182

You need to monitor resources that are distributed over different projects in Google Cloud Platform. You want to consolidate reporting under the same Stackdriver Monitoring dashboard. What should you do?

- A. Use Shared VPC to connect all projects, and link Stackdriver to one of the projects.
- B. For each project, create a Stackdriver account
- C. In each project, create a service account for that project and grant it the role of Stackdriver Account Editor in all other projects.
- D. Configure a single Stackdriver account, and link all projects to the same account.
- E. Configure a single Stackdriver account for one of the projects
- F. In Stackdriver, create a Group and add the other project names as criteria for that Group.

Answer: C

Explanation:

When you initially click on Monitoring (Stackdriver Monitoring) it creates a workspace (a stackdriver account) linked to the ACTIVE (CURRENT) Project from which it was clicked.

Now if you change the project and again click onto Monitoring it would create another workspace (a stackdriver account) linked to the changed ACTIVE (CURRENT) Project, we don't want this as this would not consolidate our results into a single dashboard (workspace/stackdriver account).

If you have accidentally created two different workspaces merge them under Monitoring > Settings > Merge Workspaces > MERGE.

If we have only one workspace and two projects we can simply add other GCP Project under Monitoring > Settings > GCP Projects > Add GCP Projects.
<https://cloud.google.com/monitoring/settings/multiple-projects>
Nothing about groups <https://cloud.google.com/monitoring/settings?hl=en>

NEW QUESTION 186

You are managing a project for the Business Intelligence (BI) department in your company. A data pipeline ingests data into BigQuery via streaming. You want the users in the BI department to be able to run the custom SQL queries against the latest data in BigQuery. What should you do?

- A. Create a Data Studio dashboard that uses the related BigQuery tables as a source and give the BI team view access to the Data Studio dashboard.
- B. Create a Service Account for the BI team and distribute a new private key to each member of the BI team.
- C. Use Cloud Scheduler to schedule a batch Dataflow job to copy the data from BigQuery to the BI team's internal data warehouse.
- D. Assign the IAM role of BigQuery User to a Google Group that contains the members of the BI team.

Answer: D

Explanation:

When applied to a dataset, this role provides the ability to read the dataset's metadata and list tables in the dataset. When applied to a project, this role also provides the ability to run jobs, including queries, within the project. A member with this role can enumerate their own jobs, cancel their own jobs, and enumerate datasets within a project. Additionally, allows the creation of new datasets within the project; the creator is granted the BigQuery Data Owner role (roles/bigquery.dataOwner) on these new datasets.
<https://cloud.google.com/bigquery/docs/access-control>

NEW QUESTION 187

You need to reduce GCP service costs for a division of your company using the fewest possible steps. You need to turn off all configured services in an existing GCP project. What should you do?

- A. * 1. Verify that you are assigned the Project Owners IAM role for this project.* 2. Locate the project in the GCP console, click Shut down and then enter the project ID.
- B. * 1. Verify that you are assigned the Project Owners IAM role for this project.* 2. Switch to the project in the GCP console, locate the resources and delete them.
- C. * 1. Verify that you are assigned the Organizational Administrator IAM role for this project.* 2. Locate the project in the GCP console, enter the project ID and then click Shut down.
- D. * 1. Verify that you are assigned the Organizational Administrators IAM role for this project.* 2. Switch to the project in the GCP console, locate the resources and delete them.

Answer: A

Explanation:

<https://cloud.google.com/run/docs/tutorials/gcloud> <https://cloud.google.com/resource-manager/docs/creating-managing-projects>
https://cloud.google.com/iam/docs/understanding-roles#primitive_roles
You can shut down projects using the Cloud Console. When you shut down a project, this immediately happens: All billing and traffic serving stops, You lose access to the project, The owners of the project will be notified and can stop the deletion within 30 days, The project will be scheduled to be deleted after 30 days. However, some resources may be deleted much earlier.

NEW QUESTION 189

Your company is moving from an on-premises environment to Google Cloud Platform (GCP). You have multiple development teams that use Cassandra environments as backend databases. They all need a development environment that is isolated from other Cassandra instances. You want to move to GCP quickly and with minimal support effort. What should you do?

- A. * 1. Build an instruction guide to install Cassandra on GCP.* 2. Make the instruction guide accessible to your developers.
- B. * 1. Advise your developers to go to Cloud Marketplace.* 2. Ask the developers to launch a Cassandra image for their development work.
- C. * 1. Build a Cassandra Compute Engine instance and take a snapshot of it.* 2. Use the snapshot to create instances for your developers.
- D. * 1. Build a Cassandra Compute Engine instance and take a snapshot of it.* 2. Upload the snapshot to Cloud Storage and make it accessible to your developers.* 3. Build instructions to create a Compute Engine instance from the snapshot so that developers can do it themselves.

Answer: B

Explanation:

<https://medium.com/google-cloud/how-to-deploy-cassandra-and-connect-on-google-cloud-platform-with-a-few->
<https://cloud.google.com/blog/products/databases/open-source-cassandra-now-managed-on-google-cloud> <https://cloud.google.com/marketplace>
You can deploy Cassandra as a Service, called Astra, on the Google Cloud Marketplace. Not only do you get a unified bill for all GCP services, your Developers can now create Cassandra clusters on Google Cloud in minutes and build applications with Cassandra as a database as a service without the operational overhead of managing Cassandra

NEW QUESTION 191

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 healthcheck 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 192

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual Associate-Cloud-Engineer Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the Associate-Cloud-Engineer Product From:

<https://www.2passeasy.com/dumps/Associate-Cloud-Engineer/>

Money Back Guarantee

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