

Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer



NEW QUESTION 1

You are working for a hospital that stores its medical images in an on-premises data room. The hospital wants to use Cloud Storage for archival storage of these images. The hospital wants an automated process to upload any new medical images to Cloud Storage. You need to design and implement a solution. What should you do?

- A. Deploy a Dataflow job from the batch template "Datastore to Cloud Storage" Schedule the batch job on the desired interval
- B. In the Cloud Console, go to Cloud Storage Upload the relevant images to the appropriate bucket
- C. Create a script that uses the gsutil command line interface to synchronize the on-premises storage with Cloud Storage Schedule the script as a cron job
- D. Create a Pub/Sub topic, and enable a Cloud Storage trigger for the Pub/Sub topic
- E. Create an application that sends all medical images to the Pub/Sub topic

Answer: C

Explanation:

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

NEW QUESTION 2

You need to create a copy of a custom Compute Engine virtual machine (VM) to facilitate an expected increase in application traffic due to a business acquisition. What should you do?

- A. Create a Compute Engine snapshot of your base VM
- B. Create your images from that snapshot.
- C. Create a Compute Engine snapshot of your base VM
- D. Create your instances from that snapshot.
- E. Create a custom Compute Engine image from a snapshot
- F. Create your images from that image.
- G. Create a custom Compute Engine image from a snapshot
- H. Create your instances from that image.

Answer: D

Explanation:

A custom image belongs only to your project. To create an instance with a custom image, you must first have a custom image.

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

- A. Create a single custom VPC with 2 subnets
- B. Create each subnet in a different region and with a different CIDR range.
- C. Create a single custom VPC with 2 subnets
- D. Create each subnet in the same region and with the same CIDR range.
- E. Create 2 custom VPCs, each with a single subnet
- F. Create each subnet in a different region and with a different CIDR range.
- G. Create 2 custom VPCs, each with a single subnet
- 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 5

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

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

Answer: C

Explanation:

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

NEW QUESTION 6

You need to select and configure compute resources for a set of batch processing jobs. These jobs take around 2 hours to complete and are run nightly. You want to minimize service costs. What should you do?

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

Answer: C

Explanation:

If your apps are fault-tolerant and can withstand possible instance preemptions, then preemptible instances can reduce your Compute Engine costs significantly. For example, batch processing jobs can run on preemptible instances. If some of those instances stop during processing, the job slows but does not completely stop. Preemptible instances complete your batch processing tasks without placing additional workload on your existing instances and without requiring you to pay full price for additional normal instances.
<https://cloud.google.com/compute/docs/instances/preemptible>

NEW QUESTION 7

You are analyzing Google Cloud Platform service costs from three separate projects. You want to use this information to create service cost estimates by service type, daily and monthly, for the next six months using standard query syntax. What should you do?

- A. Export your bill to a Cloud Storage bucket, and then import into Cloud Bigtable for analysis.
- B. Export your bill to a Cloud Storage bucket, and then import into Google Sheets for analysis.
- C. Export your transactions to a local file, and perform analysis with a desktop tool.
- D. Export your bill to a BigQuery dataset, and then write time window-based SQL queries for analysis.

Answer: D

Explanation:

"...we recommend that you enable Cloud Billing data export to BigQuery at the same time that you create a Cloud Billing account. "
<https://cloud.google.com/billing/docs/how-to/export-data-bigquery>
<https://medium.com/google-cloud/analyzing-google-cloud-billing-data-with-big-query-30bae1c2aae4>

NEW QUESTION 8

You are building an application that stores relational data from users. Users across the globe will use this application. Your CTO is concerned about the scaling requirements because the size of the user base is unknown. You need to implement a database solution that can scale with your user growth with minimum configuration changes. Which storage solution should you use?

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

Answer: B

Explanation:

Cloud Spanner is a relational database and is highly scalable. Cloud Spanner is a highly scalable, enterprise-grade, globally-distributed, and strongly consistent database service built for the cloud specifically to combine the benefits of relational database structure with a non-relational horizontal scale. This combination delivers high-performance transactions and strong consistency across rows, regions, and continents with an industry-leading 99.999% availability SLA, no planned downtime, and enterprise-grade security
Ref: <https://cloud.google.com/spanner>
Graphical user interface, application, Teams Description automatically generated

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

NEW QUESTION 9

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.

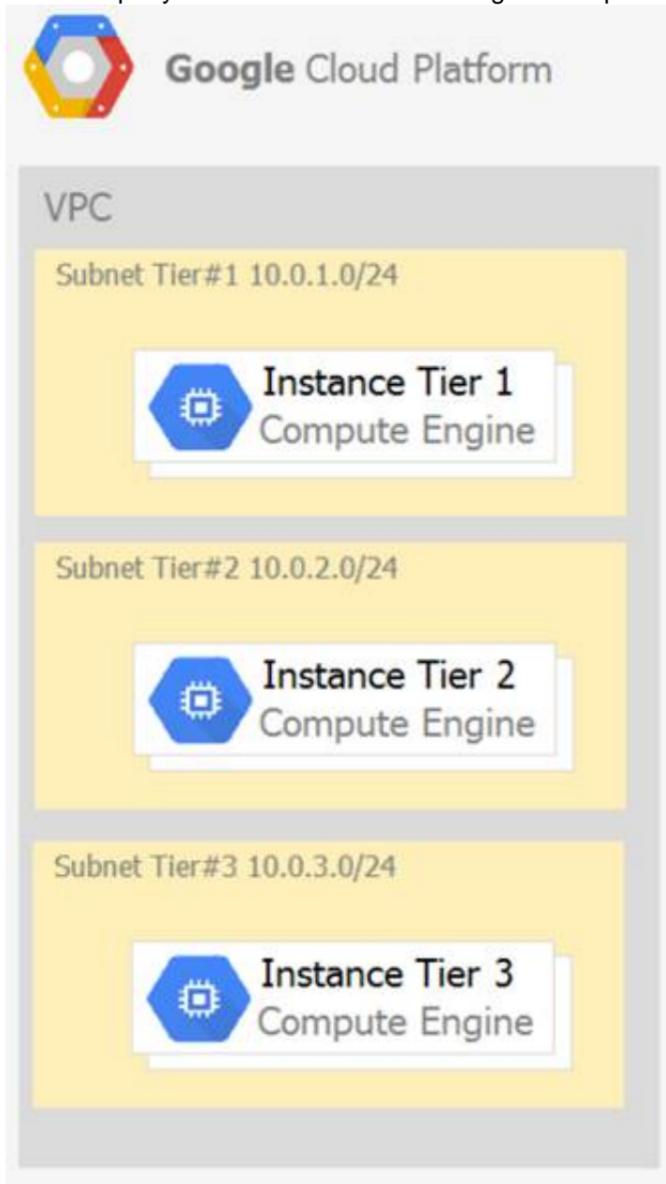
Answer: D

Explanation:

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

NEW QUESTION 10

Your company 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 10

You are hosting an application on bare-metal servers in your own data center. The application needs access to Cloud Storage. However, security policies prevent the servers hosting the application from having public IP addresses or access to the internet. You want to follow Google-recommended practices to provide the application with access to Cloud Storage. What should you do?

- A. 1. Use nslookup to get the IP address for storage.googleapis.com.2. Negotiate with the security team to be able to give a public IP address to the servers.3. Only allow egress traffic from those servers to the IP addresses for storage.googleapis.com.
- B. 1. Using Cloud VPN, create a VPN tunnel to a Virtual Private Cloud (VPC) in Google Cloud Platform (GCP).2. In this VPC, create a Compute Engine instance and install the Squid proxy server on this instance.3. Configure your servers to use that instance as a proxy to access Cloud Storage.
- C. 1. Use Migrate for Compute Engine (formerly known as Velostrata) to migrate those servers to Compute Engine.2. Create an internal load balancer (ILB) that uses storage.googleapis.com as backend.3. Configure your new instances to use this ILB as proxy.
- D. 1. Using Cloud VPN or Interconnect, create a tunnel to a VPC in GCP.2. Use Cloud Router to create a custom route advertisement for 199.36.153.4/30. Announce that network to your on-premises network through the VPN tunnel.3. In your on-premises network, configure your DNS server to resolve *.googleapis.com as a CNAME to restricted.googleapis.com.

Answer: D

Explanation:

Our requirement is to follow Google recommended practices to achieve the end result. Configuring Private Google Access for On-Premises Hosts is best achieved by VPN/Interconnect + Advertise Routes + Use restricted Google IP Range.

- Using Cloud VPN or Interconnect, create a tunnel to a VPC in GCP
- Using Cloud Router to create a custom route advertisement for 199.36.153.4/30. Announce that network to your on-premises network through the VPN tunnel.
- In your on-premises network, configure your DNS server to resolve *.googleapis.com as a CNAME to restricted.googleapis.com is the right answer right, and it is what Google recommends.

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

- You must configure routes so that Google API traffic is forwarded through your Cloud VPN or Cloud Interconnect connection, firewall rules on your on-premises firewall to allow the outgoing traffic, and DNS so that traffic to Google APIs resolves to the IP range youve added to your routes.

- You can use Cloud Router Custom Route Advertisement to announce the Restricted Google APIs IP addresses through Cloud Router to your on-premises network.

The Restricted Google APIs IP range is 199.36.153.4/30. While this is technically a public IP range, Google does not announce it publicly. This IP range is only accessible to hosts that can reach your Google Cloud projects through internal IP ranges, such as through a Cloud VPN or Cloud Interconnect connection. Without having a public IP address or access to the internet, the only way you could connect to cloud storage is if you have an internal route to it.

- So Negotiate with the security team to be able to give public IP addresses to the servers is not right.

Following Google recommended practices is synonymous with using Googles services (Not quite, but it is at least for the exam !!).

- So In this VPC, create a Compute Engine instance and install the Squid proxy server on this instance is not right.
- Migrating the VM to Compute Engine is a bit drastic when Google says it is perfectly fine to have Hybrid Connectivity architectures

<https://cloud.google.com/hybrid-connectivity>.

- So,
- Use Migrate for Compute Engine (formerly known as Velostrata) to migrate these servers to Compute Engine is not right.

NEW QUESTION 14

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 projec
- 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 19

You have a number of applications that have bursty workloads and are heavily dependent on topics to decouple publishing systems from consuming systems. Your company would like to go serverless to enable developers to focus on writing code without worrying about infrastructure. Your solution architect has already identified Cloud Pub/Sub as a suitable alternative for decoupling systems. You have been asked to identify a suitable GCP Serverless service that is easy to use

with Cloud Pub/Sub. You want the ability to scale down to zero when there is no traffic in order to minimize costs. You want to follow Google recommended practices. What should you suggest?

- A. Cloud Run for Anthos
- B. Cloud Run
- C. App Engine Standard
- D. Cloud Functions.

Answer: D

Explanation:

Cloud Functions is Google Cloud's event-driven serverless compute platform that lets you run your code locally or in the cloud without having to provision servers. Cloud Functions scales up or down, so you pay only for compute resources you use. Cloud Functions have excellent integration with Cloud Pub/Sub, lets you scale down to zero and is recommended by Google as the ideal serverless platform to use when dependent on Cloud Pub/Sub. "If you're building a simple API (a small set of functions to be accessed via HTTP or Cloud Pub/Sub), we recommend using Cloud Functions." Ref: <https://cloud.google.com/serverless-options>

NEW QUESTION 22

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

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

Answer: B

Explanation:

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

NEW QUESTION 25

You deployed an LDAP server on Compute Engine that is reachable via TLS through port 636 using UDP. You want to make sure it is reachable by clients over that port. What should you do?

- A. Add the network tag allow-udp-636 to the VM instance running the LDAP server.
- B. Create a route called allow-udp-636 and set the next hop to be the VM instance running the LDAP server.
- C. Add a network tag of your choice to the instance
- D. Create a firewall rule to allow ingress on UDP port 636 for that network tag.
- E. Add a network tag of your choice to the instance running the LDAP server
- F. Create a firewall rule to allow egress on UDP port 636 for that network tag.

Answer: C

Explanation:

A tag is simply a character string added to a tags field in a resource, such as Compute Engine virtual machine (VM) instances or instance templates. A tag is not a separate resource, so you cannot create it separately. All resources with that string are considered to have that tag. Tags enable you to make firewall rules and routes applicable to specific VM instances.

NEW QUESTION 28

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 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 35

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

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

Answer: B

Explanation:

The key to understand the requirement is : "The objects are written once and accessed frequently for 30 days" Standard Storage

Standard Storage is best for data that is frequently accessed ("hot" data) and/or stored for only brief periods of time.

Archive Storage

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

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

NEW QUESTION 36

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

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

Answer: D

Explanation:

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

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

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

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

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

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

NEW QUESTION 41

You are storing sensitive information in a Cloud Storage bucket. For legal reasons, you need to be able to record all requests that read any of the stored data. You want to make sure you comply with these requirements. What should you do?

- A. Enable the Identity Aware Proxy API on the project.
- B. Scan the bucket using the Data Loss Prevention API.
- C. Allow only a single Service Account access to read the data.
- D. Enable Data Access audit logs for the Cloud Storage API.

Answer: D

Explanation:

Logged information Within Cloud Audit Logs, there are two types of logs: Admin Activity logs: Entries for operations that modify the configuration or metadata of a project, bucket, or object. Data Access logs: Entries for operations that modify objects or read a project, bucket, or object. There are several sub-types of data access logs: ADMIN_READ: Entries for operations that read the configuration or metadata of a project, bucket, or object. DATA_READ: Entries for operations that read an object. DATA_WRITE: Entries for operations that create or modify an object. <https://cloud.google.com/storage/docs/audit-logs#types>

NEW QUESTION 42

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 43

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 48

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 49

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 50

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?

- Create a cron job that runs on a scheduled basis to review stackdriver monitoring metrics, and then resize the Spanner instance accordingly.
- Create a Stackdriver alerting policy to send an alert to oncall SRE emails when Cloud Spanner CPU exceeds the threshold
- SREs would scale resources up or down accordingly.
- Create a Stackdriver alerting policy to send an alert to Google Cloud Support email when Cloud Spanner CPU exceeds your threshold
- Google support would scale resources up or down accordingly.
- Create a Stackdriver alerting policy to send an alert to webhook when Cloud Spanner CPU is over or under your threshold
- 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 55

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 57

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 60

You created several resources in multiple Google Cloud projects. All projects are linked to different billing accounts. To better estimate future charges, you want to have a single visual representation of all costs incurred. You want to include new cost data as soon as possible. What should you do?

- A. Configure Billing Data Export to BigQuery and visualize the data in Data Studio.
- B. Visit the Cost Table page to get a CSV export and visualize it using Data Studio.
- C. Fill all resources in the Pricing Calculator to get an estimate of the monthly cost.
- D. Use the Reports view in the Cloud Billing Console to view the desired cost information.

Answer: A

Explanation:

<https://cloud.google.com/billing/docs/how-to/export-data-bigquery> "Cloud Billing export to BigQuery enables you to export detailed Google Cloud billing data (such as usage, cost estimates, and pricing data) automatically throughout the day to a BigQuery dataset that you specify."

NEW QUESTION 65

You want to find out when users were added to Cloud Spanner Identity Access Management (IAM) roles on your Google Cloud Platform (GCP) project. What should you do in the GCP Console?

- A. Open the Cloud Spanner console to review configurations.
- B. Open the IAM & admin console to review IAM policies for Cloud Spanner roles.
- C. Go to the Stackdriver Monitoring console and review information for Cloud Spanner.
- D. Go to the Stackdriver Logging console, review admin activity logs, and filter them for Cloud Spanner IAM roles.

Answer: D

Explanation:

<https://cloud.google.com/monitoring/audit-logging>

NEW QUESTION 70

You are given a project with a single virtual private cloud (VPC) and a single subnetwork in the us-central1 region. There is a Compute Engine instance hosting an application in this subnetwork. You need to deploy a new instance in the same project in the europe-west1 region. This new instance needs access to the application. You want to follow Google-recommended practices. What should you do?

- A. 1. Create a subnetwork in the same VPC, in europe-west1.2. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.
- B. 1. Create a VPC and a subnetwork in europe-west1.2. Expose the application with an internal load balancer.3. Create the new instance in the new subnetwork and use the load balancer's address as the endpoint.
- C. 1. Create a subnetwork in the same VPC, in europe-west1.2. Use Cloud VPN to connect the two subnetworks.3. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.
- D. 1. Create a VPC and a subnetwork in europe-west1.2. Peer the 2 VPCs.3. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.

Answer: C

Explanation:

> Given that the new instance wants to access the application on the existing compute engine instance, these applications seem to be related so they should be within the same VPC. It is possible to have them in different VPCs and peer the VPCs but this is a lot of additional work and we can simplify this by choosing the option below (which is the answer)

* 1. Create a subnet in the same VPC, in europe-west1.

* 2. Create the new instance in the new subnet and use the first instance subnets private address as the endpoint. is the right answer.

> We can create another subnet in the same VPC and this subnet is located in europe-west1. We can then spin up a new instance in this subnet. We also have to set up a firewall rule to allow communication between the two subnets. All instances in the two subnets with the same VPC can communicate through the internal IP Address

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

NEW QUESTION 72

You are building a product on top of Google Kubernetes Engine (GKE). You have a single GKE cluster. For each of your customers, a Pod is running in that cluster, and your customers can run arbitrary code inside their Pod. You want to maximize the isolation between your customers' Pods. What should you do?

- A. Use Binary Authorization and whitelist only the container images used by your customers' Pods.
- B. Use the Container Analysis API to detect vulnerabilities in the containers used by your customers' Pods.
- C. Create a GKE node pool with a sandbox type configured to gvisor
- D. Add the parameter runtimeClassName: gvisor to the specification of your customers' Pods.
- E. Use the cos_containerd image for your GKE node
- F. Add a nodeSelector with the value cloud.google.com/gke-os-distribution: cos_containerd to the specification of your customers' Pods.

Answer: C

NEW QUESTION 75

You need to update a deployment in Deployment Manager without any resource downtime in the deployment. Which command should you use?

- A. `gcloud deployment-manager deployments create --config <deployment-config-path>`
- B. `gcloud deployment-manager deployments update --config <deployment-config-path>`
- C. `gcloud deployment-manager resources create --config <deployment-config-path>`
- D. `gcloud deployment-manager resources update --config <deployment-config-path>`

Answer: B

NEW QUESTION 76

You need to provide a cost estimate for a Kubernetes cluster using the GCP pricing calculator for Kubernetes. Your workload requires high IOPs, and you will also be using disk snapshots. You start by entering the number of nodes, average hours, and average days. What should you do next?

- A. Fill in local SS
- B. Fill in persistent disk storage and snapshot storage.
- C. Fill in local SS
- D. Add estimated cost for cluster management.
- E. Select Add GPU
- F. Fill in persistent disk storage and snapshot storage.
- G. Select Add GPU
- H. Add estimated cost for cluster management.

Answer: A

Explanation:

<https://cloud.google.com/compute/docs/disks/local-ssd>

NEW QUESTION 78

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 80

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 82

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, individual
- 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, individual
- 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 85

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 90

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

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

Answer: C

Explanation:

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

gcloud config configurations list lists existing named configurations

gcloud config configurations activate activates an existing named configuration

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

NEW QUESTION 92

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

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

Answer: D

NEW QUESTION 93

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 97

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 100

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 setting
- 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 103

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 104

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 109

You created a Google Cloud Platform project with an App Engine application inside the project. You initially configured the application to be served from the us-central region. Now you want the application to be served from the asia-northeast1 region. What should you do?

- A. Change the default region property setting in the existing GCP project to asia-northeast1.
- B. Change the region property setting in the existing App Engine application from us-central to asia-northeast1.
- C. Create a second App Engine application in the existing GCP project and specify asia-northeast1 as the region to serve your application.
- D. Create a new GCP project and create an App Engine application inside this new project.
- E. Specify asia-northeast1 as the region to serve your application.

Answer: D

Explanation:

<https://cloud.google.com/appengine/docs/flexible/managing-projects-apps-billing#:~:text=Each%20Cloud%20p> Two App engine can't be running on the same project: you can check this easy diagram for more info:

https://cloud.google.com/appengine/docs/standard/an-overview-of-app-engine#components_of_an_application

And you can't change location after setting it for your app Engine. <https://cloud.google.com/appengine/docs/standard/locations>

App Engine is regional and you cannot change an apps region after you set it. Therefore, the only way to have an app run in another region is by creating a new project and targeting the app engine to run in the required region (asia-northeast1 in our case).

Ref: <https://cloud.google.com/appengine/docs/locations>

NEW QUESTION 114

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 116

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 117

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 120

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 122

You have a Compute Engine instance hosting a production application. You want to receive an email if the instance consumes more than 90% of its CPU resources for more than 15 minutes. You want to use Google services. What should you do?

- A. * 1. Create a consumer Gmail account.* 2. Write a script that monitors the CPU usage.* 3. When the CPU usage exceeds the threshold, have that script send an email using the Gmail account and smtp.gmail.com on port 25 as SMTP server.
- B. * 1. Create a Stackdriver Workspace, and associate your Google Cloud Platform (GCP) project with it.* 2.Create an Alerting Policy in Stackdriver that uses the threshold as a trigger condition
- C. 3.Configure your email address in the notification channel.
- D. * 1. Create a Stackdriver Workspace, and associate your GCP project with it.* 2.Write a script that monitors the CPU usage and sends it as a custom metric to Stackdriver
- E. 3.Create an uptime check for the instance in Stackdriver.
- F. * 1. In Stackdriver Logging, create a logs-based metric to extract the CPU usage by using this regular expression: CPU Usage: ([0-9] {1,3})%* 2.In Stackdriver Monitoring, create an Alerting Policy based on this metric
- G. 3.Configure your email address in the notification channel.

Answer: B

Explanation:

Specifying conditions for alerting policies This page describes how to specify conditions for alerting policies. The conditions for an alerting policy define what is monitored and when to trigger an alert. For example, suppose you want to define an alerting policy that emails you if the CPU utilization of a Compute Engine VM instance is above 80% for more than 3 minutes. You use the conditions dialog to specify that you want to monitor the CPU utilization of a Compute Engine VM instance, and that you want an alerting policy to trigger when that utilization is above 80% for 3 minutes. <https://cloud.google.com/monitoring/alerts/ui-conditions-ga>
<https://cloud.google.com/monitoring/alerts/using-alerting-ui> <https://cloud.google.com/monitoring/support/notification-options>

NEW QUESTION 126

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 VM's 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 127

You have successfully created a development environment in a project for an application. This application uses Compute Engine and Cloud SQL. Now, you need to create a production environment for this application.

The security team has forbidden the existence of network routes between these 2 environments, and asks you to follow Google-recommended practices. What should you do?

- A. Create a new project, enable the Compute Engine and Cloud SQL APIs in that project, and replicate the setup you have created in the development environment.
- B. Create a new production subnet in the existing VPC and a new production Cloud SQL instance in your existing project, and deploy your application using those resources.
- C. Create a new project, modify your existing VPC to be a Shared VPC, share that VPC with your new project, and replicate the setup you have in the development environment in that new project, in the Shared VPC.
- D. Ask the security team to grant you the Project Editor role in an existing production project used by another division of your company.
- E. Once they grant you that role, replicate the setup you have in the development environment in that project.

Answer: A

Explanation:

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

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

NEW QUESTION 132

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: D

Explanation:

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

NEW QUESTION 133

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 to 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 138

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

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

Answer: D

Explanation:

Keyword need to look for

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

NEW QUESTION 140

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 145

You are using Deployment Manager to create a Google Kubernetes Engine cluster. Using the same Deployment Manager deployment, you also want to create a DaemonSet in the kube-system namespace of the cluster. You want a solution that uses the fewest possible services. What should you do?

- A. Add the cluster's API as a new Type Provider in Deployment Manager, and use the new type to create the DaemonSet.
- B. Use the Deployment Manager Runtime Configurator to create a new Config resource that contains the DaemonSet definition.
- C. With Deployment Manager, create a Compute Engine instance with a startup script that uses kubectl to create the DaemonSet.
- D. In the cluster's definition in Deployment Manager, add a metadata that has kube-system as key and the DaemonSet manifest as value.

Answer: A

Explanation:

Adding an API as a type provider

This page describes how to add an API to Google Cloud Deployment Manager as a type provider. To learn more about types and type providers, read the Types overview documentation.

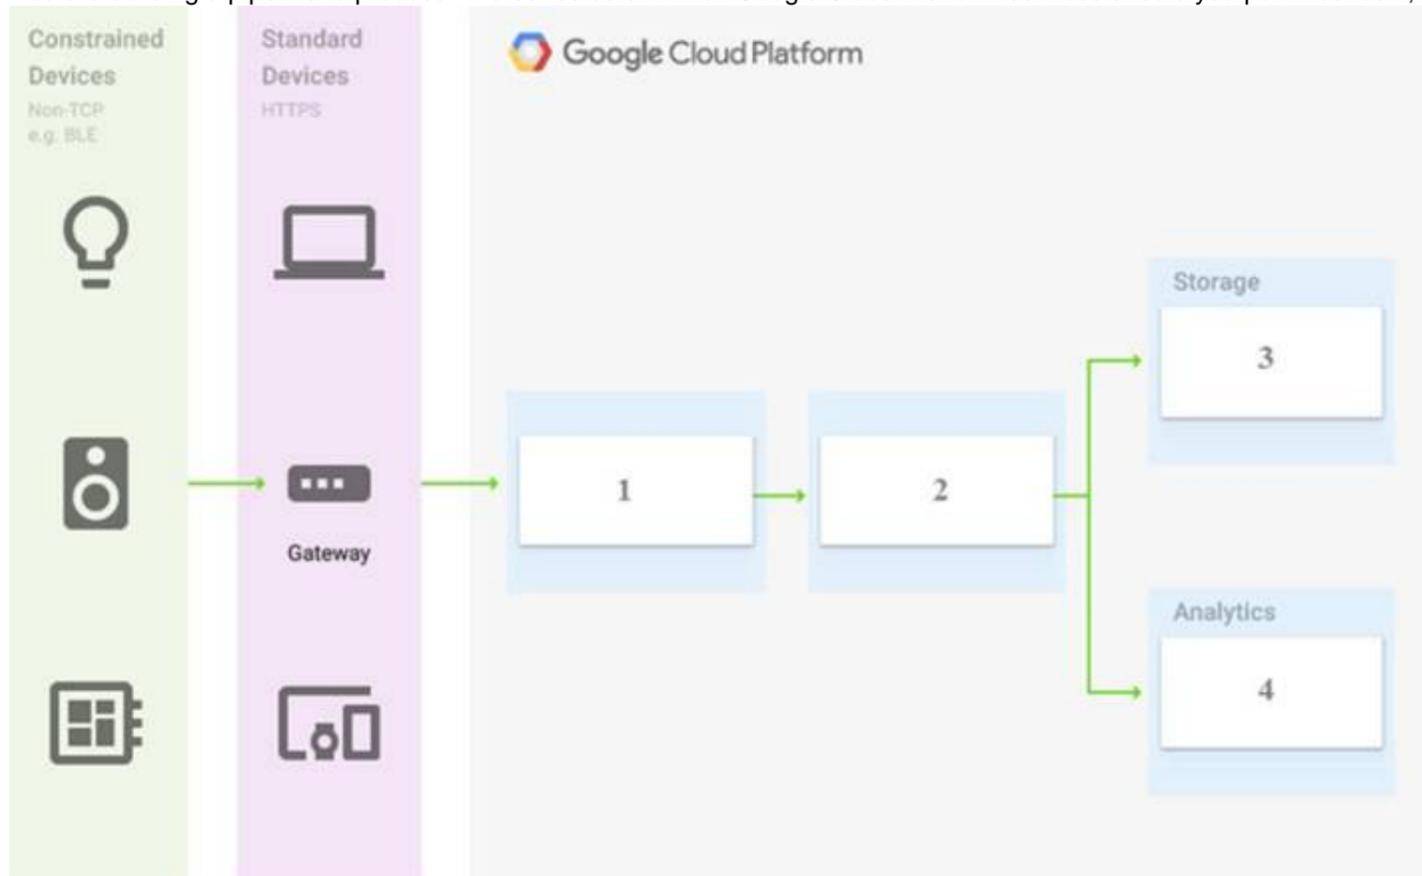
A type provider exposes all of the resources of a third-party API to Deployment Manager as base types that you can use in your configurations. These types must be directly served by a RESTful API that supports Create, Read, Update, and Delete (CRUD).

If you want to use an API that is not automatically provided by Google with Deployment Manager, you must add the API as a type provider.

<https://cloud.google.com/deployment-manager/docs/configuration/type-providers/creating-type-provider>

NEW QUESTION 148

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 153

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

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

Answer: A

NEW QUESTION 154

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

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

Answer: A

NEW QUESTION 157

You have a development project with appropriate IAM roles defined. You are creating a production project and want to have the same IAM roles on the new project, using the fewest possible steps. What should you do?

- A. Use `gcloud iam roles copy` and specify the production project as the destination project.
- B. Use `gcloud iam roles copy` and specify your organization as the destination organization.
- C. In the Google Cloud Platform Console, use the 'create role from role' functionality.
- D. In the Google Cloud Platform Console, use the 'create role' functionality and select all applicable permissions.

Answer: A

NEW QUESTION 158

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 163

You are configuring service accounts for an application that spans multiple projects. Virtual machines (VMs) running in the web-applications project need access to BigQuery datasets in `crm-databases-proj`. You want to follow Google-recommended practices to give access to the service account in the web-applications project. What should you do?

- A. Give "project owner" for web-applications appropriate roles to `crm-databases-proj`
- B. Give "project owner" role to `crm-databases-proj` and the web-applications project.
- C. Give "project owner" role to `crm-databases-proj` and `bigquery.dataViewer` role to web-applications.
- D. Give `bigquery.dataViewer` role to `crm-databases-proj` and appropriate roles to web-applications.

Answer: D

NEW QUESTION 167

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 168

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

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

Answer: A

Explanation:

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

NEW QUESTION 170

Your development team needs a new Jenkins server for their project. You need to deploy the server using the fewest steps possible. What should you do?

- A. Download and deploy the Jenkins Java WAR to App Engine Standard.
- B. Create a new Compute Engine instance and install Jenkins through the command line interface.
- C. Create a Kubernetes cluster on Compute Engine and create a deployment with the Jenkins Docker image.
- D. Use GCP Marketplace to launch the Jenkins solution.

Answer: D

NEW QUESTION 174

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 177

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

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

Answer: A

Explanation:

Create an instance template for the instances so VMs have same specs. Set the 'Automatic Restart' to on to VM automatically restarts upon crash. Set the 'On-

host maintenance' to Migrate VM instance. This will take care of VM during maintenance window. It will migrate VM instance making it highly available Add the instance template to an instance group so instances can be managed.

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

Enabling automatic restart ensures that compute engine instances are automatically restarted when they crash. And Enabling Migrate VM Instance enables live migrates i.e. compute instances are migrated during system maintenance and remain running during the migration.

Automatic Restart If your instance is set to terminate when there is a maintenance event, or if your instance crashes because of an underlying hardware issue, you can set up Compute Engine to automatically restart the instance by setting the automaticRestart field to true. This setting does not apply if the instance is taken offline through a user action, such as calling sudo shutdown, or during a zone

outage.Ref: <https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options#autorestart>

Enabling the Migrate VM Instance option migrates your instance away from an infrastructure maintenance event, and your instance remains running during the migration. Your instance might experience a short period of decreased performance, although generally, most instances should not notice any difference. This is ideal for instances that require constant uptime and can tolerate a short period of decreased

performance.Ref: https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options#live_

NEW QUESTION 179

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?

- Deploy the container on Cloud Run.
- Deploy the container on Cloud Run on GKE.
- Deploy the container on App Engine Flexible.
- 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 182

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?

- Create an HTTP load balancer with a backend configuration that references an existing instance group.Set the health check to healthy (HTTP).
- 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.
- Create a managed instance grou
- Set the Autohealing health check to healthy (HTTP).
- Create a managed instance grou
- Verify that the autoscaling setting is on.

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 187

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

- Create a service account with just the permissions to access files in the bucke
- Create a JSON key for the service accoun
- Execute the command `gsutil signurl -m 1h gs:///*`.
- Create a service account with just the permissions to access files in the bucke
- Create a JSON key for the service accoun
- Execute the command `gsutil signurl -d 1h gs:///**`.
- Create a service account with just the permissions to access files in the bucke
- Create a JSON key for the service accoun
- Execute the command `gsutil signurl -p 60m gs:///`.
- Create a JSON key for the Default Compute Engine Service Accoun
- Execute the command `gsutil signurl -t 60m gs:///***`

Answer: B

Explanation:

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

NEW QUESTION 190

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 191

Several employees at your company have been creating projects with Cloud Platform and paying for it with their personal credit cards, which the company reimburses. The company wants to centralize all these projects under a single, new billing account. What should you do?

- A. Contact cloud-billing@google.com with your bank account details and request a corporate billing account for your company.
- B. Create a ticket with Google Support and wait for their call to share your credit card details over the phone.
- C. In the Google Platform Console, go to the Resource Manager and move all projects to the root Organization.
- D. In the Google Cloud Platform Console, create a new billing account and set up a payment method.

Answer: D

Explanation:

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

NEW QUESTION 196

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 role stage 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 200

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 202

You have designed a solution on Google Cloud Platform (GCP) that uses multiple GCP products. Your company has asked you to estimate the costs of the solution. You need to provide estimates for the monthly total cost. What should you do?

- A. For each GCP product in the solution, review the pricing details on the products pricing pag
- B. Use the pricing calculator to total the monthly costs for each GCP product.
- C. For each GCP product in the solution, review the pricing details on the products pricing pag
- D. Create a Google Sheet that summarizes the expected monthly costs for each product.
- E. Provision the solution on GC
- F. Leave the solution provisioned for 1 wee
- G. Navigate to the Billing Report page in the Google Cloud Platform Consol
- H. Multiply the 1 week cost to determine the monthly costs.
- I. Provision the solution on GC
- J. Leave the solution provisioned for 1 wee
- K. Use Stackdriver to determine the provisioned and used resource amount
- L. Multiply the 1 week cost to determine the monthly costs.

Answer: A

Explanation:

You can use the Google Cloud Pricing Calculator to total the estimated monthly costs for each GCP product. You dont incur any charges for doing so.
Ref: <https://cloud.google.com/products/calculator>

NEW QUESTION 206

Your Dataproc cluster runs in a single Virtual Private Cloud (VPC) network in a single subnet with range 172.16.20.128/25. There are no private IP addresses available in the VPC network. You want to add new VMs to communicate with your cluster using the minimum number of steps. What should you do?

- A. Modify the existing subnet range to 172.16.20.0/24.
- B. Create a new Secondary IP Range in the VPC and configure the VMs to use that range.
- C. Create a new VPC network for the VM
- D. Enable VPC Peering between the VMs' VPC network and the Dataproc cluster VPC network.
- E. Create a new VPC network for the VMs with a subnet of 172.32.0.0/16. Enable VPC network Peering between the Dataproc VPC network and the VMs VPC network
- F. Configure a custom Route exchange.

Answer: A

Explanation:

/25:
CIDR to IP Range Result
CIDR Range 172.16.20.128/25 Netmask 255.255.255.128
Wildcard Bits 0.0.0.127
First IP 172.16.20.128
First IP (Decimal) 2886734976 Last IP 172.16.20.255
Last IP (Decimal) 2886735103 Total Host 128
CIDR 172.16.20.128/25
/24:
CIDR to IP Range Result
CIDR Range 172.16.20.128/24 Netmask 255.255.255.0
Wildcard Bits 0.0.0.255
First IP 172.16.20.0
First IP (Decimal) 2886734848 Last IP 172.16.20.255
Last IP (Decimal) 2886735103 Total Host 256
CIDR 172.16.20.128/24

NEW QUESTION 208

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

Answer: A

Explanation:

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

NEW QUESTION 213

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 project
- 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 an another workspace(a stackdriver account) linked to the changed ACTIVE(CURRENT) Project, we don't want this as this would not consolidate our result into a single dashboard(workspace/stackdriver account).

If you have accidentally created two diff 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 215

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 219

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 220

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 225

.....

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](#)