Google

Exam Questions Professional-Cloud-Architect

Google Certified Professional - Cloud Architect (GCP)

NEW QUESTION 1
- (Topic 1)
For this question, refer to the Mountkirk Games case study
Mountkirk Games needs to create a repeatable and configurable mechanism for deploying isolated application environments. Developers and testers can access each other's environments and resources, but they cannot access staging or production resources. The staging environment needs access to some services from production.
What should you do to isolate development environments from staging and production?

A. Create a project for development and test and another for staging and production.
B. Create a network for development and test and another for staging and production.
C. Create one subnetwork for development and another for staging and production.
D. Create one project for development, a second for staging and a third for production.

Answer: D

NEW QUESTION 2
- (Topic 2)
For this question, refer to the TerramEarth case study
Your development team has created a structured API to retrieve vehicle data. They want to allow third parties to develop tools for dealerships that use this vehicle event data. You want to support delegated authorization against this data. What should you do?

A. Build or leverage an OAuth-compatible access control system.
B. Build SAML 2.0 SSO compatibility into your authentication system.
C. Restrict data access based on the source IP address of the partner systems.
D. Create secondary credentials for each dealer that can be given to the trusted third party.

Answer: A

Explanation:
https://cloud.google.com/appengine/docs/flexible/go/authorizing-apps
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#delegate_application_authorization_with_oauth2
Delegate application authorization with OAuth2
Cloud Platform APIs support OAuth 2.0, and scopes provide granular authorization over the methods that are supported. Cloud Platform supports both service-account and user- account OAuth, also called three-legged OAuth.
References: https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#delegate_application_authorization_with_oauth2
https://cloud.google.com/appengine/docs/flexible/go/authorizing-apps

NEW QUESTION 3
- (Topic 2)
For this question, refer to the TerramEarth case study.
The TerramEarth development team wants to create an API to meet the company's business requirements. You want the development team to focus their development effort on business value versus creating a custom framework. Which method should they use?

A. Use Google App Engine with Google Cloud Endpoint
B. Focus on an API for dealers and partners.
C. Use Google App Engine with a JAX-RS Jersey Java-based framewor
D. Focus on an API for the public.
E. Use Google App Engine with the Swagger (open API Specification) framewor
F. Focus on an API for the public.
G. Use Google Container Engine with a Django Python containe
H. Focus on an API for the public.
I. Use Google Container Engine with a Tomcat container with the Swagger (Open API Specification) framewor
J. Focus on an API for dealers and partners.

Answer: A

Explanation:
https://cloud.google.com/endpoints/docs/openapi/about-cloud- endpoints?hl=en_US&_ga=2.21787131.-1712523161.1522785064
https://cloud.google.com/endpoints/docs/openapi/architecture-overview https://cloud.google.com/storage/docs/gsutil/commands/test
Develop, deploy, protect and monitor your APIs with Google Cloud Endpoints. Using an Open API Specification or one of our API frameworks, Cloud Endpoints gives you the tools you need for every phase of API development.
From scenario: Business Requirements
Decrease unplanned vehicle downtime to less than 1 week, without increasing the cost of carrying surplus inventory
Support the dealer network with more data on how their customers use their equipment to better position new products and services
Have the ability to partner with different companies – especially with seed and fertilizer suppliers in the fast-growing agricultural business – to create compelling joint offerings for their customers.
Reference: https://cloud.google.com/certification/guides/cloud-architect/casestudy- terramearth

NEW QUESTION 4
- (Topic 2)
For this question, refer to the TerramEarth case study
You analyzed TerramEarth's business requirement to reduce downtime, and found that they can achieve a majority of time saving by reducing customers' wait time for parts You decided to focus on reduction of the 3 weeks aggregate reporting time Which modifications to the company's processes should you recommend?

A. Migrate from CSV to binary format, migrate from FTP to SFTP transport, and develop machine learning analysis of metrics.
B. Migrate from FTP to streaming transport, migrate from CSV to binary format, and develop machine learning analysis of metrics.
C. Increase fleet cellular connectivity to 80%, migrate from FTP to streaming transport, and develop machine learning analysis of metrics.
D. Migrate from FTP to SFTP transport, develop machine learning analysis of metrics, and increase dealer local inventory by a fixed factor.

Answer: C

Explanation:
The Avro binary format is the preferred format for loading compressed data. Avro data is faster to load because the data can be read in parallel, even when the data blocks are compressed.
Cloud Storage supports streaming transfers with the gsutil tool or boto library, based on HTTP chunked transfer encoding. Streaming data lets you stream data to and from your Cloud Storage account as soon as it becomes available without requiring that the data be first saved to a separate file. Streaming transfers are useful if you have a process that generates data and you do not want to buffer it locally before uploading it, or if you want to send the result from a computational pipeline directly into Cloud Storage.
References: https://cloud.google.com/storage/docs/streaming https://cloud.google.com/bigquery/docs/loading-data

NEW QUESTION 5
- (Topic 3)
For this question, refer to the JencoMart case study.
JencoMart has built a version of their application on Google Cloud Platform that serves traffic to Asia. You want to measure success against their business and technical goals.
Which metrics should you track?

A. Error rates for requests from Asia
B. Latency difference between US and Asia
C. Total visits, error rates, and latency from Asia
D. Total visits and average latency for users in Asia
E. The number of character sets present in the database

Answer: D

NEW QUESTION 6
- (Topic 3)
For this question, refer to the JencoMart case study.
The JencoMart security team requires that all Google Cloud Platform infrastructure is deployed using a least privilege model with separation of duties for administration between production and development resources. What Google domain and project structure should you recommend?

A. Create two G Suite accounts to manage users: one for development/test/staging andone for productio
B. Each account should contain one project for every application.
C. Create two G Suite accounts to manage users: one with a single project for all development applications and one with a single project for all production applications.
D. Create a single G Suite account to manage users with each stage of each application in its own project.
E. Create a single G Suite account to manage users with one project for the development/test/staging environment and one project for the production environment.

Answer: D

Explanation:
Note: The principle of least privilege and separation of duties are concepts that, although semantically different, are intrinsically related from the standpoint of security. The intent behind both is to prevent people from having higher privilege levels than they actually need
? Principle of Least Privilege: Users should only have the least amount of privileges required to perform their job and no more. This reduces authorization exploitation by limiting access to resources such as targets, jobs, or monitoring templates for which they are not authorized.
? Separation of Duties: Beyond limiting user privilege level, you also limit user duties, or the specific jobs they can perform. No user should be given responsibility for more than one related function. This limits the ability of a user to perform a malicious action and then cover up that action.
References: https://cloud.google.com/kms/docs/separation-of-duties

NEW QUESTION 7
- (Topic 3)
For this question, refer to the JencoMart case study
A few days after JencoMart migrates the user credentials database to Google Cloud Platform and shuts down the old server, the new database server stops responding to SSH connections. It is still serving database requests to the application servers correctly. What three steps should you take to diagnose the problem? Choose 3 answers

A. Delete the virtual machine (VM) and disks and create a new one.
B. Delete the instance, attach the disk to a new VM, and investigate.
C. Take a snapshot of the disk and connect to a new machine to investigate.
D. Check inbound firewall rules for the network the machine is connected to.
E. Connect the machine to another network with very simple firewall rules and investigate.
F. Print the Serial Console output for the instance for troubleshooting, activate the interactive console, and investigate.

Answer: CDF

Explanation:
https://cloud.google.com/compute/docs/troubleshooting/troubleshooting-ssh
D: Handling "Unable to connect on port 22" error message Possible causes include:
There is no firewall rule allowing SSH access on the port. SSH access on port 22 is enabled on all Compute Engine instances by default. If you have disabled access, SSH from the Browser will not work. If you run sshd on a port other than 22, you need to enable the access to that port with a custom firewall rule.
The firewall rule allowing SSH access is enabled, but is not configured to allow connections from GCP Console services. Source IP addresses for browser-based SSH sessions are dynamically allocated by GCP Console and can vary from session to session.
References: https://cloud.google.com/compute/docs/ssh-in-browser https://cloud.google.com/compute/docs/ssh-in-browser

NEW QUESTION 8
- (Topic 4)
For this question, refer to the Dress4Win case study.
As part of their new application experience, Dress4Wm allows customers to upload images of themselves. The customer has exclusive control over who may view

these images. Customers should be able to upload images with minimal latency and also be shown their images quickly on the main application page when they log in. Which configuration should Dress4Win use?

A. Store image files in a Google Cloud Storage bucke
B. Use Google Cloud Datastore to maintain metadata that maps each customer's ID and their image files.
C. Store image files in a Google Cloud Storage bucke
D. Add custom metadata to the uploaded images in Cloud Storage that contains the customer's unique ID.
E. Use a distributed file system to store customers' image
F. As storage needs increase, add more persistent disks and/or node
G. Assign each customer a unique ID, which sets each file's owner attribute, ensuring privacy of images.
H. Use a distributed file system to store customers' image
I. As storage needs increase, add more persistent disks and/or node
J. Use a Google Cloud SQL database to maintain metadata that maps each customer's ID to their image files.

Answer: A

NEW QUESTION 9
- (Topic 4)
The current Dress4win system architecture has high latency to some customers because it is located in one data center.
As of a future evaluation and optimizing for performance in the cloud, Dresss4win wants to distribute it's system architecture to multiple locations when Google cloud platform. Which approach should they use?

A. Use regional managed instance groups and a global load balancer to increase performance because theregional managed instance group can grow instances in each region separately based on traffic.
B. Use a global load balancer with a set of virtual machines that forward the requests to a closer group ofvirtual machines managed by your operations team.
C. Use regional managed instance groups and a global load balancer to increase reliability by providingautomatic failover between zones in different regions.
D. Use a global load balancer with a set of virtual machines that forward the requests to a closer group ofvirtual machines as part of a separate managed instance groups.

Answer: A

NEW QUESTION 10
- (Topic 5)
You are responsible for the Google Cloud environment in your company Multiple departments need access to their own projects and the members within each department will have the same project responsibilities You want to structure your Google Cloud environment for minimal maintenance and maximum overview of 1AM permissions as each department's projects start and end You want to follow Google-recommended practices What should you do?

A. Create a Google Group per department and add all department members to their respective groups Create a folder per departmentand grant the respective group the required 1AM permissions at the folder level Add the projects under the respective folders
B. Grant all department members the required 1AM permissions for their respective projects
C. Create a Google Group per department and add all department members to theirrespective groups Grant each group the required I AM permissions for their respective projects
D. Create a folder per department and grant the respective members of the department the required 1AM permissions at the folder leve
E. Structure all projects for each department under the respective folders

Answer: A

Explanation:
This option follows the Google-recommended practices for structuring a Google Cloud environment for minimal maintenance and maximum overview of IAM permissions. By creating a Google Group per department and adding all department members to their respective groups, you can simplify user management and avoid granting IAM permissions to individual users. By creating a folder per department and granting the respective group the required IAM permissions at the folder level, you can enforce consistent policies across all projects within each department and avoid granting IAM permissions at the project level. By adding the projects under the respective folders, you can organize your resources hierarchically and leverage inheritance of IAM policies from folders to projects. The other options are not optimal for this scenario, because they either require granting IAM permissions to individual users (B, C), or do not use Google Groups to manage users (D). References:
? https://cloud.google.com/architecture/framework/system-design
? https://cloud.google.com/architecture/identity/best-practices-for-planning
? https://cloud.google.com/resource-manager/docs/creating-managing-folders

NEW QUESTION 10
- (Topic 5)
Your company wants you to build a highly reliable web application with a few public APIs as the backend. You don’t expect a lot of user traffic, but traffic could spike occasionally.
You want to leverage Cloud Load Balancing, and the solution must be cost-effective for users. What should you do?

A. Store static content such as HTML and images in Cloud CD
B. Host the APIs on App Engine and store the user data in Cloud SQL.
C. Store static content such as HTML and images in a Cloud Storage bucke
D. Host the APIs on a zonal Google Kubernetes Engine cluster with worker nodes in multiple zones, and save the user data in Cloud Spanner.
E. Store static content such as HTML and images in Cloud CD
F. Use Cloud Run to host the APIs and save the user data in Cloud SQL.
G. Store static content such as HTML and images in a Cloud Storage bucke
H. Use Cloud Functions to host the APIs and save the user data in Firestore.

Answer: D

Explanation:
https://cloud.google.com/load-balancing/docs/https/setting-up-https-serverless#gcloud:-cloud-functions https://cloud.google.com/blog/products/networking/better-load-balancing-for-app-engine-cloud-run-and-functions

NEW QUESTION 12
- (Topic 5)
Your company is planning to upload several important files to Cloud Storage. After the upload is completed, they want to verify that the upload content is identical to what they have on- premises. You want to minimize the cost and effort of performing this check. What should you do?
A.
1) Use gsutil -m to upload all the files to Cloud Storage.
2) Use gsutil cp to download the uploaded files
3) Use Linux diff to compare the content of the files
B.
1) Use gsutil -m to upload all the files to Cloud Storage.
2) Develop a custom Java application that computes CRC32C hashes
3) Use gsutil ls -L gs://[YOUR_BUCKET_NAME] to collect CRC32C hashes of the uploaded files
4) Compare the hashes
C.
1) Use Linux shasum to compute a digest of files you want to upload
2) Use gsutil -m to upload all the files to the Cloud Storage
3) Use gsutil cp to download the uploaded files
4) Use Linux shasum to compute a digest of the downloaded files 5.Compre the hashes
D.
1) Use gsutil -m to upload all the files to Cloud Storage.
2) Use gsutil hash -c FILE_NAME to generate CRC32C hashes of all on-premises files 3)Use gsutil ls -L gs://[YOUR_BUCKET_NAME] to collect CRC32C hashes of the uploaded files
4)Compare the hashes

A.

Answer: D

Explanation:
https://cloud.google.com/storage/docs/gsutil/commands/hash

NEW QUESTION 14
- (Topic 5)
You have a Python web application with many dependencies that requires 0.1 CPU cores and 128 MB of memory to operate in production. You want to monitor and maximize machine utilization. You also to reliably deploy new versions of the application. Which set of steps should you take?

A. Perform the following:1) Create a managed instance group with f1-micro type machines.2) Use a startup script to clone the repository, check out the production branch, install the dependencies, and start the Python app.3) Restart the instances to automatically deploy new production releases.
B. Perform the following:1) Create a managed instance group with n1-standard-1 type machines.2) Build a Compute Engine image from the production branch that contains all of the dependencies andautomatically starts the Python app.3) Rebuild the Compute Engine image, and update the instance template to deploy new productionreleases.
C. Perform the following:1) Create a Kubernetes Engine cluster with n1-standard-1 type machines.2) Build a Docker image from the production branch with all of the dependencies, and tag it with theversion number.3) Create a Kubernetes Deployment with the imagePullPolicy set to “IfNotPresent” in the stagingnamespace, and then promote it to the production namespace after testing.
D. Perform the following:1) Create a Kubernetes Engine (GKE) cluster with n1-standard-4 type machines.2) Build a Docker image from the master branch will all of the dependencies, and tag it with “latest”.3) Create a Kubernetes Deployment in the default namespace with the imagePullPolicy set to “Always”.Restart the pods to automatically deploy new production releases.

Answer: D

Explanation:
https://cloud.google.com/compute/docs/instance-templates

NEW QUESTION 15
- (Topic 5)
You need to ensure reliability for your application and operations by supporting reliable task a scheduling for compute on GCP. Leveraging Google best practices, what should you do?

A. Using the Cron service provided by App Engine, publishing messages directly to a message-processing utility service running on Compute Engine instances.
B. Using the Cron service provided by App Engine, publish messages to a Cloud Pub/Sub topi
C. Subscribe to that topic using a message-processing utility service running on Compute Engine instances.
D. Using the Cron service provided by Google Kubernetes Engine (GKE), publish messages directly to amessage-processing utility service running on Compute Engine instances.
E. Using the Cron service provided by GKE, publish messages to a Cloud Pub/Sub topi
F. Subscribe to that topic using a message-processing utility service running on Compute Engine instances.

Answer: B

Explanation:
https://cloud.google.com/solutions/reliable-task-scheduling-compute-engine

NEW QUESTION 16
- (Topic 5)
Your company is developing a web-based application. You need to make sure that production deployments are linked to source code commits and are fully auditable. What should you do?

A. Make sure a developer is tagging the code commit with the date and time of commit
B. Make sure a developer is adding a comment to the commit that links to the deployment.
C. Make the container tag match the source code commit hash.

D. Make sure the developer is tagging the commits with :latest

Answer: C

Explanation:
From: https://cloud.google.com/architecture/best-practices-for-building- containers
Under: Tagging using the Git commit hash (bottom of page almost)
"In this case, a common way of handling version numbers is to use the Git commit SHA-1 hash (or a short version of it) as the version number. By design, the Git commit hash is immutable and references a specific version of your software.
You can use this commit hash as a version number for your software, but also as a tag for the Docker image built from this specific version of your software. Doing so makes Docker images traceable: because in this case the image tag is immutable, you instantly know which specific version of your software is running inside a given container."

NEW QUESTION 19
- (Topic 5)
Your company has a Google Cloud project that uses BigQuery for data warehousing on a pay-per-use basis. You want to monitor queries in real time to discover the most costly queries and which users spend the most. What should you do?
A.
* 1. Create a Cloud Logging sink to export BigQuery data access logs to Cloud Storage.
* 2. Develop a Dataflow pipeline to compute the cost of queries split by users.
B.
* 1. Create a Cloud Logging sink to export BigQuery data access logs to BigQuery.
* 2. Perform a BigQuery query on the generated table to extract the information you need.
C.
* 1. Activate billing export into BigQuery.
* 2. Perform a BigQuery query on the billing table to extract the information you need.
D.
* 1. In the BigQuery dataset that contains all the tables to be queried, add a label for each user that can launch a query.
* 2. Open the Billing page of the project.
* 3. Select Reports.
* 4. Select BigQuery as the product and filter by the user you want to check.

A.

Answer: C

Explanation:
https://cloud.google.com/blog/products/data-analytics/taking-a-practical-approach-to-bigquery-cost-monitoring

NEW QUESTION 21
- (Topic 5)
Your company is planning to perform a lift and shift migration of their Linux RHEL 6.5+ virtual machines. The virtual machines are running in an on-premises VMware environment. You want to migrate them to Compute Engine following Google- recommended practices. What should you do?

A. * 1. Define a migration plan based on the list of the applications and their dependencies.* 2. Migrate all virtual machines into Compute Engine individually with Migrate for Compute Engine.
B. * 1. Perform an assessment of virtual machines running in the current VMware environment.* 2. Create images of all disk
C. Import disks on Compute Engine.* 3. Create standard virtual machines where the boot disks are the ones you have imported.
D. * 1. Perform an assessment of virtual machines running in the current VMware environment.* 2. Define a migration plan, prepare a Migrate for Compute Engine migration RunBook, and execute the migration.
E. * 1. Perform an assessment of virtual machines running in the current VMware environment.* 2.Install a third-party agent on all selected virtual machine
F. 3.Migrate all virtual machines into Compute Engine.

Answer: C

Explanation:
The framework illustrated in the preceding diagram has four phases:
•Assess. In this phase, you assess your source environment, assess the workloads that you want to migrate to Google Cloud, and assess which VMs support each workload.
•Plan. In this phase, you create the basic infrastructure for Migrate for Compute Engine, such as provisioning the resource hierarchy and setting up network access.
•Deploy. In this phase, you migrate the VMs from the source environment to Compute Engine.
•Optimize. In this phase, you begin to take advantage of the cloud technologies and capabilities.
Reference: https://cloud.google.com/architecture/migrating-vms-migrate-for-compute-engine-getting-started

NEW QUESTION 22
- (Topic 5)
You are tasked with building an online analytical processing (OLAP) marketing analytics and reporting tool.
This requires a relational database that can operate on hundreds of terabytes of data. What is the Google recommended tool for such applications?

A. Cloud Spanner, because it is globally distributed
B. Cloud SQL, because it is a fully managed relational database
C. Cloud Firestore, because it offers real-time synchronization across devices
D. BigQuery, because it is designed for large-scale processing of tabular data

Answer: A

Explanation:
Reference: https://cloud.google.com/files/BigQueryTechnicalWP.pdf

NEW QUESTION 26
- (Topic 5)
You are deploying a PHP App Engine Standard service with SQL as the backend. You want to minimize the number of queries to the database.
What should you do?

A. Set the memcache service level to dedicate
B. Create a key from the hash of the query, and returndatabase values from memcache before issuing a query to Cloud SQL.
C. Set the memcache service level to dedicate
D. Create a cron task that runs every minute to populate the cache with keys containing query results.
E. Set the memcache service level to share
F. Create a cron task that runs every minute to save all expected queries to a key called “cached-queries”.
G. Set the memcache service level to share
H. Create a key called “cached-queries”, and return databasevalues from the key before using a query to Cloud SQL.

Answer: A

Explanation:
https://cloud.google.com/appengine/docs/standard/php/memcache/using

NEW QUESTION 27
- (Topic 5)
Your operations team currently stores 10 TB of data m an object storage service from a third-party provider. They want to move this data to a Cloud Storage bucket as quickly as possible, following Google-recommended practices. They want to minimize the cost of this data migration. When approach should they use?

A. Use the gsutil mv command lo move the data
B. Use the Storage Transfer Service to move the data
C. Download the data to a Transfer Appliance and ship it to Google
D. Download the data to the on-premises data center and upload it to the Cloud Storage bucket

Answer: B

Explanation:
https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#transfer-options
https://cloud.google.com/storage-transfer-service

NEW QUESTION 32
- (Topic 5)
Your customer wants to do resilience testing of their authentication layer. This consists of a regional managed instance group serving a public REST API that reads from and writes to a Cloud SQL instance.
What should you do?

A. Engage with a security company to run web scrapes that look your users’ authentication data om malicious websites and notify you if any if found.
B. Deploy intrusion detection software to your virtual machines to detect and log unauthorized access.
C. Schedule a disaster simulation exercise during which you can shut off all VMs in a zone to see how your application behaves.
D. Configure a red replica for your Cloud SQL instance in a different zone than the master,and then manually trigger a failover while monitoring KPIs for our REST API.

Answer: C

NEW QUESTION 35
- (Topic 5)
You deploy your custom Java application to Google App Engine. It fails to deploy and gives you the following stack trace.

```
java.lang.SecurityException: SHA1 digest error for
com/Altostrat/CloakedServlet.class
     at com.google.appengine.runtime.Request.process
-d36f818a24b8cf1d (Request.java)
     at
sun.security.util.ManifestEntryVerifier.verify
(ManifestEntryVerifier.java:210)
     at java.util.jar.JarVerifier.processEntry
(JarVerifier.java:218)
     at java.util.jar.JarVerifier.update
(JarVerifier.java:205)
     at
java.util.jar.JarVerifiersVerifierStream.read
(JarVerifier.java:428)
     at sun.misc.Resource.getBytes
(Resource.java:124)
     at java.net.URL.ClassLoader.defineClass
(URLClassLoader.java:273)
     at sun.reflect.GeneratedMethodAccessor5.invoke
(Unknown Source)
     at
sun.reflect.DelegatingMethodAccessorImpl.invoke
(DelegatingMethodAccessorImpl.java:43)
     at java.lang.reflect.Method.invoke
(Method.java:616)
     at java.lang.ClassLoader.loadClass
(ClassLoader.java:266)
```

What should you do?

A. Upload missing JAR files and redeploy your application.
B. Digitally sign all of your JAR files and redeploy your application
C. Recompile the CLoakedServlet class using and MD5 hash instead of SHA1

Answer: B

NEW QUESTION 39
- (Topic 5)
One of the developers on your team deployed their application in Google Container Engine with the Dockerfile below. They report that their application deployments are taking too long.

```
FROM ubuntu:16.04

COPY . /src

RUN apt-get update && apt-get install -y python python-pip

RUN pip install -r requirements.txt
```

You want to optimize this Dockerfile for faster deployment times without adversely affecting the app's functionality.
Which two actions should you take? Choose 2 answers.

A. Remove Python after running pip.
B. Remove dependencies from requirements.txt.
C. Use a slimmed-down base image like Alpine linux.
D. Use larger machine types for your Google Container Engine node pools.
E. Copy the source after the package dependencies (Python and pip) are installed.

Answer: CE

Explanation:
The speed of deployment can be changed by limiting the size of the uploaded app, limiting the complexity of the build necessary in the Dockerfile, if present, and by ensuring a fast and reliable internet connection.
Note: Alpine Linux is built around musl libc and busybox. This makes it smaller and more resource efficient than traditional GNU/Linux distributions. A container requires no more
than 8 MB and a minimal installation to disk requires around 130 MB of storage. Not only do you get a fully-fledged Linux environment but a large selection of packages from the repository.
References: https://groups.google.com/forum/#!topic/google-appengine/hZMEkmmObDU https://www.alpinelinux.org/about/

NEW QUESTION 42
- (Topic 5)
Your operations team has asked you to help diagnose a performance issue in a production application that runs on Compute Engine. The application is dropping requests that reach it when under heavy load. The process list for affected instances shows a single application process that is consuming all available CPU, and autoscaling has reached the upper limit of instances. There is no abnormal load on any other related systems, including the database. You want to allow production traffic to be served again as quickly as possible. Which action should you recommend?

A. Change the autoscaling metric to agent.googleapis.com/memory/percent_used.
B. Restart the affected instances on a staggered schedule.
C. SSH to each instance and restart the application process.
D. Increase the maximum number of instances in the autoscaling group.

Answer: D

Explanation:
Reference: https://cloud.google.com/blog/products/sap-google-cloud/best-practices-for- sap-app-server- autoscaling-on-google-cloud

NEW QUESTION 46
- (Topic 5)
Your company runs several databases on a single MySQL instance. They need to take backups of a specific database at regular intervals. The backup activity needs to complete as quickly as possible and cannot be allowed to impact disk performance. How should you configure the storage?

A. Configure a cron job to use the gcloud tool to take regular backups using persistent disk snapshots.
B. Mount a Local SSD volume as the backup locatio
C. After the backup is complete, use gsutil to move the backup to Google Cloud Storage.
D. Use gcsfuse to mount a Google Cloud Storage bucket as a volume directly on the instance and write backups to the mounted location using mysqldump
E. Mount additional persistent disk volumes onto each virtual machine (VM) instance in a RAID10 array and use LVM to create snapshots to send to Cloud Storage.

Answer: B

Explanation:
https://cloud.google.com/compute/docs/instances/sql-server/best-practices

NEW QUESTION 49
- (Topic 5)
You need to reduce the number of unplanned rollbacks of erroneous production deployments in your company's web hosting platform. Improvement to the QA/Test processes accomplished an 80% reduction. Which additional two approaches can you take to further reduce the rollbacks? Choose 2 answers

A. Introduce a green-blue deployment model.
B. Replace the QA environment with canary releases.
C. Fragment the monolithic platform into microservices.
D. Reduce the platform's dependency on relational database systems.
E. Replace the platform's relational database systems with a NoSQL database.

Answer: AC

NEW QUESTION 50
- (Topic 5)
You are working with a data warehousing team that performs data analysis. The team needs to process data from external partners, but the data contains personally identifiable information (PII). You need to process and store the data without storing any of the Pll data. What should you do?

A. Create a Dataflow pipeline to retrieve the data from the external source
B. As part of the pipeline use the Cloud Data Loss Prevention (Cloud DLP) API to remove any Pll data Store the result in BigQuery
C. Create a Dataflow pipeline to retrieve the data from the external source
D. As part of the pipeline store all non-PII data in BigQuery and store all Pll data in a Cloud Storage bucket that has a retention policy set.
E. Ask the external partners to upload an data on Cloud Storage Configure Bucket Lock for the bucket Create a Dataflow pipeline to read the data from the bucket As part of the pipeline, use the Cloud Data Loss Prevention (Cloud DIP) API to remove any Pll data Store the result in BigQuery
F. Ask the external partners to import ail data in your BigQuery dataset Create a dataflow pipeline to copy the data into a new table As part of the Dataflow bucket skip all data in columns that have Pll data

Answer: A

Explanation:
Create a Dataflow pipeline to retrieve the data from the external sources, he
did not specify the way he is going to create it, it might be a pub/sub or external table or whatever.

NEW QUESTION 54
- (Topic 5)
You are running a cluster on Kubernetes Engine to serve a web application. Users are reporting that a specific part of the application is not responding anymore. You notice that all pods of your deployment keep restarting after 2 seconds. The application writes logs to standard output. You want to inspect the logs to find the cause of the issue. Which approach can you take?

A. Review the Stackdriver logs for each Compute Engine instance that is serving as a node in the cluster.
B. Review the Stackdriver logs for the specific Kubernetes Engine container that is serving the unresponsive part of the application.
C. Connect to the cluster using gcloud credentials and connect to a container in one of the pods to read the logs.
D. Review the Serial Port logs for each Compute Engine instance that is serving as a node in the cluster.

Answer: B

NEW QUESTION 58
- (Topic 5)
Your development team has installed a new Linux kernel module on the batch servers in Google Compute Engine (GCE) virtual machines (VMs) to speed up the nightly batch process. Two days after the installation, 50% of web application deployed in the same
nightly batch run. You want to collect details on the failure to pass back to the development team. Which three actions should you take? Choose 3 answers

A. Use Stackdriver Logging to search for the module log entries.
B. Read the debug GCE Activity log using the API or Cloud Console.
C. Use gcloud or Cloud Console to connect to the serial console and observe the logs.
D. Identify whether a live migration event of the failed server occurred, using in the activity log.
E. Adjust the Google Stackdriver timeline to match the failure time, and observe the batch server metrics.
F. Export a debug VM into an image, and run the image on a local server where kernel log messages will be displayed on the native screen.

Answer: ACE

Explanation:
https://www.flexera.com/blog/cloud/2013/12/google-compute-engine-live-migration-passes- the-test/
"With live migration, the virtual machines are moved without any downtime or noticeable service degradation"

NEW QUESTION 62
- (Topic 5)
Your company has an application running as a Deployment in a Google Kubernetes Engine (GKE) cluster When releasing new versions of the application via a rolling deployment, the team has been causing outages The root cause of the outages is misconfigurations with parameters that are only used in production You want to put preventive measures for this in the platform to prevent outages What should you do?

A. Configure liveness and readiness probes in the Pod specification
B. Configure an uptime alert in Cloud Monitoring
C. Create a Scheduled Task to check whether the application is available
D. Configure health checks on the managed instance group

Answer: D

Explanation:
This option can help prevent outages caused by misconfigurations with parameters that are only used in production. Liveness and readiness probes are mechanisms to check the health and availability of the Pods and containers in a GKE cluster. Liveness probes determine if a container is still running, and if not, restart it. Readiness probes determine if a container is ready to serve requests, and if not, remove it from the load balancer. By configuring liveness and readiness probes in the Pod specification, you can ensure that your application can handle traffic and recover from failures gracefully during a rolling update. The other options are not optimal for this scenario, because they either do not prevent outages, but only alert or monitor them (B, C), or do not apply to GKE clusters, but to Compute Engine instances (D). References:
? https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps
? https://cloud.google.com/blog/products/containers-kubernetes/kubernetes-best- practices-setting-up-health-checks-with-readiness-and-liveness-probes

NEW QUESTION 65
- (Topic 5)
Your company has a Google Cloud project that uses BigQuery for data warehousing. The VPN tunnel between the on-premises environment and Google Cloud is configured with Cloud VPN. Your security team wants to avoid data exfiltration by malicious insiders, compromised code, and accidental oversharing. What should you do?

A. Configure VPC Service Controls and configure Private Google Access for on-premises hosts.
B. Create a service account, grant the BigQuery JobUser role and Storage Object Viewer role to the service account, and remove all other Identity and Access Management (1AM) access from the project.
C. Configure Private Google Access.
D. Configure Private Service Connect.

Answer: A

NEW QUESTION 69
- (Topic 5)
You are configuring the cloud network architecture for a newly created project m Google Cloud that will host applications in Compote Engine Compute Engine virtual machine instances will be created in two different subnets (sub-a and sub-b) within a single region
• Instances in sub-a win have public IP addresses
• Instances in sub-b will have only private IP addresses
To download updated packages, instances must connect to a public repository outside the boundaries of Google Cloud You need to allow sub-b to access the external repository. What should you do?

A. Enable Private Google Access on sub-b
B. Configure Cloud NAT and select sub b m the NAT mapping section
C. Configure a bastion host instance in sub a to connect to instances in sub-b
D. Enable Identity Aware Proxy for TCP forwarding for instances in sub-b

Answer: B

Explanation:
? Cloud NAT (network address translation) lets Google Cloud virtual machine (VM) instances without external IP addresses and private Google Kubernetes Engine (GKE) clusters send outbound packets to the internet and receive any corresponding established inbound response packets1. By configuring Cloud NAT and selecting sub-b in the NAT mapping section, you can allow instances in sub-b to access the external repository without exposing them to the internet1.

NEW QUESTION 73

- (Topic 5)
Your company has announced that they will be outsourcing operations functions. You want to allow developers to easily stage new versions of a cloud-based application in the production environment and allow the outsourced operations team to autonomously promote staged versions to production. You want to minimize the operational overhead of the solution. Which Google Cloud product should you migrate to?

A. App Engine
B. GKE On-Prem
C. Compute Engine
D. Google Kubernetes Engine

Answer: A

Explanation:
Reference: https://cloud.google.com/security/compliance/eba-outsourcing-mapping-gcp

NEW QUESTION 75
- (Topic 5)
Your company is running a stateless application on a Compute Engine instance. The application is used
heavily during regular business hours and lightly outside of business hours. Users are reporting that the application is slow during peak hours. You need to optimize the application's performance. What should you do?

A. Create a snapshot of the existing dis
B. Create an instance template from the snapshot.Create anautoscaled managed instance group from the instance template.
C. Create a snapshot of the existing dis
D. Create a custom image from the snapsho
E. Create an autoscaled managed instance group from the custom image.
F. Create a custom image from the existing dis
G. Create an instance template from the custom imag
H. Create an autoscaled managed instance group from the instance template.
I. Create an instance template from the existing dis
J. Create a custom image from the instance template.Create an autoscaled managed instance group from the custom image.

Answer: B

Explanation:
https://cloud.google.com/compute/docs/instance-templates/create-instance- templates

NEW QUESTION 78
- (Topic 5)
Your customer support tool logs all email and chat conversations to Cloud Bigtable for retention and analysis.
What is the recommended approach for sanitizing this data of personally identifiable information or payment
card information before initial storage?

A. Hash all data using SHA256
B. Encrypt all data using elliptic curve cryptography
C. De-identify the data with the Cloud Data Loss Prevention API
D. Use regular expressions to find and redact phone numbers, email addresses, and credit card numbers

Answer: A

Explanation:
Reference: https://cloud.google.com/solutions/pci-dss-compliance-ingcp#

NEW QUESTION 80
- (Topic 5)
You have created several preemptible Linux virtual machine instances using Google Compute Engine. You want to properly shut down your application before the virtual machines are preempted. What should you do?

A. Create a shutdown script named k99.shutdown in the /etc/rc.6.d/ directory.
B. Create a shutdown script registered as a xinetd service in Linux and configure a Stackdnver endpoint check to call the service.
C. Create a shutdown script and use it as the value for a new metadata entry with the key shutdown-script in the Cloud Platform Console when you create the new virtual machine instance.
D. Create a shutdown script, registered as a xinetd service in Linux, and use the gcloud compute instances add-metadata command to specify the service URL as the value for a new metadata entry with the key shutdown-script-url

Answer: C

NEW QUESTION 83
- (Topic 5)
You are managing several projects on Google Cloud and need to interact on a daily basis with BigQuery, Bigtable and Kubernetes Engine using the gcloud CLI tool You are travelling a lot and work on different workstations during the week You want to avoid having to manage the gcloud CLI manually What should you do?

A. Use a package manager to install gcloud on your workstations instead of installing it manually
B. Create a Compute Engine instance and install gcloud on the instance Connect to this instance via SSH to always use the samegcloud installation when interacting with Google Cloud
C. Install gcloud on all of your workstations Run the command gcloud components auto- update on each workstation
D. Use Google Cloud Shell in the Google Cloud Console to interact with Google Cloud

Answer: D

Explanation:
This option allows you to use the gcloud CLI tool without having to install or manage it manually on different workstations. Google Cloud Shell is a browser-based command-line tool that provides you with a temporary Compute Engine virtual machine instance preloaded with the Cloud SDK, including the gcloud CLI tool. You can access Google Cloud Shell from any web browser and use it to interact with BigQuery, Bigtable and Kubernetes Engine using the gcloud CLI tool. The other options are not optimal for this scenario, because they either require installing and updating the gcloud CLI tool on multiple workstations (A, C), or creating and maintaining a Compute Engine instance for the sole purpose of using the gcloud CLI tool (B). References:
? https://cloud.google.com/shell/docs/overview
? https://cloud.google.com/sdk/gcloud/

NEW QUESTION 86
- (Topic 5)
You need to migrate Hadoop jobs for your company's Data Science team without modifying the underlying infrastructure. You want to minimize costs and infrastructure management effort. What should you do?

A. Create a Dataproc cluster using standard worker instances.
B. Create a Dataproc cluster using preemptible worker instances.
C. Manually deploy a Hadoop cluster on Compute Engine using standard instances.
D. Manually deploy a Hadoop cluster on Compute Engine using preemptible instances.

Answer: B

Explanation:
Reference: https://cloud.google.com/architecture/hadoop/hadoop-gcp-migration-jobs

NEW QUESTION 88
- (Topic 5)
You have an App Engine application that needs to be updated. You want to test the update with production traffic before replacing the current application version. What should you do?

A. Deploy the update using the Instance Group Updater to create a partial rollout, which allows for canarytesting.
B. Deploy the update as a new version in the App Engine application, and split traffic between the new and current versions.
C. Deploy the update in a new VPC, and use Google's global HTTP load balancing to split traffic between the update and current applications.
D. Deploy the update as a new App Engine application, and use Google's global HTTP load balancing to split traffic between the new and current applications.

Answer: B

Explanation:
https://cloud.google.com/appengine/docs/standard/python/splitting-traffic

NEW QUESTION 90
- (Topic 5)
You need to evaluate your team readiness for a new GCP project. You must perform the evaluation and create a skills gap plan incorporates the business goal of cost optimization. Your team has deployed two GCP projects successfully to date. What should you do?

A. Allocate budget for team trainin
B. Set a deadline for the new GCP project.
C. Allocate budget for team trainin
D. Create a roadmap for your team to achieve Google Cloud certification based on job role.
E. Allocate budget to hire skilled external consultant
F. Set a deadline for the new GCP project.
G. Allocate budget to hire skilled external consultant
H. Create a roadmap for your team to achieve Google Cloud certification based on job role.

Answer: B

Explanation:
https://services.google.com/fh/files/misc/cloud_center_of_excellence.pdf

NEW QUESTION 93
- (Topic 5)
You write a Python script to connect to Google BigQuery from a Google Compute Engine virtual machine. The script is printing errors that it cannot connect to BigQuery. What should you do to fix the script?

A. Install the latest BigQuery API client library for Python
B. Run your script on a new virtual machine with the BigQuery access scope enabled
C. Create a new service account with BigQuery access and execute your script with that user
D. Install the bq component for gccloud with the command gcloud components install bq.

Answer: B

Explanation:
The error is most like caused by the access scope issue. When create new instance, you have the default Compute engine default service account but most serves access including BigQuery is not enable. Create an instance Most access are not enabled by default You have default service account but don't have the permission (scope) you can stop the instance, edit, change scope and restart it to enable the scope access. Of course, if you Run your script on a new virtual machine with the BigQuery access scope enabled, it also works
https://cloud.google.com/compute/docs/access/service-accounts

NEW QUESTION 94

- (Topic 5)
You team needs to create a Google Kubernetes Engine (GKE) cluster to host a newly built application that requires access to third-party services on the internet. Your company does not allow any Compute Engine instance to have a public IP address on Google Cloud. You need to create a deployment strategy that adheres to these guidelines. What should you do?

A. Create a Compute Engine instance, and install a NAT Proxy on the instanc
B. Configure all workloads on GKE to pass through this proxy to access third-party services on the Internet
C. Configure the GKE cluster as a private cluster, and configure Cloud NAT Gateway for the cluster subnet
D. Configure the GKE cluster as a route-based cluste
E. Configure Private Google Access on the Virtual Private Cloud (VPC)
F. Configure the GKE cluster as a private cluste
G. Configure Private Google Access on the Virtual Private Cloud (VPC)

Answer: B

Explanation:
A Cloud NAT gateway can perform NAT for nodes and Pods in a private cluster, which is a type of VPC-native cluster. The Cloud NAT gateway must be configured to apply to at least the following subnet IP address ranges for the subnet that your cluster uses:
Subnet primary IP address range (used by nodes)
Subnet secondary IP address range used for Pods in the cluster Subnet secondary IP address range used for Services in the cluster
The simplest way to provide NAT for an entire private cluster is to configure a Cloud NAT gateway to apply to all of the cluster's subnet's IP address ranges.
https://cloud.google.com/nat/docs/overview

NEW QUESTION 98
- (Topic 5)
You company has a Kubernetes application that pulls messages from Pub/Sub and stores them in Firestore. Because the application is simple, it was deployed as a single pod. The infrastructure team has analyzed Pub/Sub metrics and discovered that the application cannot process the messages in real time. Most of them wait for minutes before being processed. You need to scale the elaboration process that is I/O-intensive. What should you do?

A. Configure a Kubernetes autoscaling based on the subscription/push_request metric.
B. Use the –enable- autoscaling flag when you create the Kubernetes cluster
C. Configure a Kubernetes autoscaling based on the subscription/num_undelivered message metric.
D. Use kubectl autoscale deployment APP_NAME –max 6 –min 2 –cpu- percent 50 to configure Kubernetes autoscaling deployment

Answer: A

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/custom-and-external-metrics#external_metrics

NEW QUESTION 99
- (Topic 5)
You have developed a non-critical update to your application that is running in a managed instance group, and have created a new instance template with the update that you want to release. To prevent any possible impact to the application, you don't want to update any running instances. You want any new instances that are created by the managed instance group to contain the new update. What should you do?

A. Start a new rolling restart operation.
B. Start a new rolling replace operation.
C. Start a new rolling updat
D. Select the Proactive update mode.
E. Start a new rolling updat
F. Select the Opportunistic update mode.

Answer: D

Explanation:
In certain scenarios, an opportunistic update is useful because you don't want to cause instability to the system if it can be avoided. For example, if you have a non- critical update that can be applied as necessary without any urgency and you have a MIG that is actively being autoscaled, perform an opportunistic update so that Compute Engine does not actively tear down your existing instances to apply the update. When resizing down, the autoscaler preferentially terminates instances with the old template as well as instances that are not yet in a RUNNING state.

NEW QUESTION 100
- (Topic 5)
Your company pushes batches of sensitive transaction data from its application server VMs to Cloud Pub/Sub for processing and storage. What is the Google-recommended way for your application to authenticate to the required Google Cloud services?

A. Ensure that VM service accounts are granted the appropriate Cloud Pub/Sub IAM roles.
B. Ensure that VM service accounts do not have access to Cloud Pub/Sub, and use VM access scopes togrant the appropriate Cloud Pub/Sub IAM roles.
C. Generate an OAuth2 access token for accessing Cloud Pub/Sub, encrypt it, and store it in Cloud Storage for access from each VM.
D. Create a gateway to Cloud Pub/Sub using a Cloud Function, and grant the Cloud Function service account the appropriate Cloud Pub/Sub IAM roles.

Answer: A

NEW QUESTION 101
- (Topic 5)
Your company is developing a new application that will allow globally distributed users to upload pictures and share them with other selected users. The application will support millions of concurrent users. You want to allow developers to focus on just building code without having to create and maintain the underlying infrastructure. Which service should you use to deploy the application?

A. App Engine

B. Cloud Endpoints
C. Compute Engine
D. Google Kubernetes Engine

Answer: A

Explanation:
Reference: https://cloud.google.com/terms/services https://cloud.google.com/appengine/docs/standard/go/how-requests-are-handled

NEW QUESTION 103
- (Topic 5)
Your architecture calls for the centralized collection of all admin activity and VM system logs within your project.
How should you collect these logs from both VMs and services?

A. All admin and VM system logs are automatically collected by Stackdriver.
B. Stackdriver automatically collects admin activity logs for most service
C. The Stackdriver Logging agentmust be installed on each instance to collect system logs.
D. Launch a custom syslogd compute instance and configure your GCP project and VMs to forward all logs to it.
E. Install the Stackdriver Logging agent on a single compute instance and let it collect all audit and access logs for your environment.

Answer: B

Explanation:
https://cloud.google.com/logging/docs/agent/default-logs

NEW QUESTION 104
- (Topic 5)
You want to enable your running Google Kubernetes Engine cluster to scale as demand for your application changes.
What should you do?

A. Add additional nodes to your Kubernetes Engine cluster using the following command:gcloud container clusters resizeCLUSTER_Name – -size 10
B. Add a tag to the instances in the cluster with the following command:gcloud compute instances add-tagsINSTANCE - -tags enable-autoscaling max-nodes-10
C. Update the existing Kubernetes Engine cluster with the following command:gcloud alpha container clustersupdate mycluster - -enable-autoscaling - -min-nodes=1 - -max-nodes=10
D. Create a new Kubernetes Engine cluster with the following command:gcloud alpha container clusterscreate mycluster - -enable-autoscaling - -min-nodes=1 - -max- nodes=10and redeploy your application

Answer: C

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler To enable autoscaling for an existing node pool, run the following command:
gcloud container clusters update [CLUSTER_NAME] --enable-autoscaling \--min-nodes 1 -- max-nodes 10 --zone [COMPUTE_ZONE] --node-pool default-pool

NEW QUESTION 108
- (Topic 5)
Your company wants to migrate their 10-TB on-premises database export into Cloud Storage You want to minimize the time it takes to complete this activity, the overall cost and database load The bandwidth between the on-premises environment and Google Cloud is 1 Gbps You want to follow Google-recommended practices What should you do?

A. Use the Data Transfer appliance to perform an offline migration
B. Use a commercial partner ETL solution to extract the data from the on-premises database and upload it into Cloud Storage
C. Develop a Dataflow job to read data directly from the database and write it into Cloud Storage
D. Compress the data and upload it with gsutii -m to enable multi-threaded copy

Answer: A

Explanation:
The Data Transfer appliance is a Google-provided hardware device that can be used to transfer large amounts of data from on-premises environments to Cloud Storage. It is suitable for scenarios where the bandwidth between the on-premises environment and Google Cloud is low or insufficient, and the data size is large. The Data Transfer appliance can minimize the time it takes to complete the migration, the overall cost and database load, by avoiding network bottlenecks and reducing bandwidth consumption. The Data Transfer appliance also encrypts the data at rest and in transit,
ensuring data security and privacy. The other options are not optimal for this scenario, because they either require a high-bandwidth network connection (B, C, D), or incur additional costs and complexity (B, C). References:
? https://cloud.google.com/data-transfer-appliance/docs/overview
? https://cloud.google.com/blog/products/storage-data-transfer/introducing-storage- transfer-service-for-on-premises-data

NEW QUESTION 109
- (Topic 5)
You are using Cloud SQL as the database backend for a large CRM deployment. You want to scale as usage increases and ensure that you don’t run out of storage, maintain 75% CPU usage cores, and keep replication lag below 60 seconds. What are the correct steps to meet your requirements?

A. 1) Enable automatic storage increase for the instance.2) Create a Stackdriver alert when CPU usage exceeds 75%, and change the instance type to reduceCPU usage.3) Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
B. 1) Enable automatic storage increase for the instance.2) Change the instance type to a 32-core machine type to keep CPU usage below 75%.3) Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
C. 1) Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on theinstance to create more space.2) Deploy memcached to reduce CPU load.3) Change the instance type to a 32-core machine type to reduce replication lag.
D. 1) Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on theinstance to create more space.2) Deploy memcached to

reduce CPU load.3) Create a Stackdriver alert for replication lag, and change the instance type to a 32-core machine type to reduce replication lag.

Answer: A

NEW QUESTION 113
- (Topic 5)
A development team at your company has created a dockerized HTTPS web application. You need to deploy the application on Google Kubernetes Engine (GKE) and make sure that the application scales automatically.
How should you deploy to GKE?

A. Use the Horizontal Pod Autoscaler and enable cluster autoscalin
B. Use an Ingressresource to loadbalance the HTTPS traffic.
C. Use the Horizontal Pod Autoscaler and enable cluster autoscaling on the Kubernetes cluste
D. Use aService resource of type LoadBalancer to load-balance the HTTPS traffic.
E. Enable autoscaling on the Compute Engine instance grou
F. Use an Ingress resource to load balance the HTTPS traffic.
G. Enable autoscaling on the Compute Engine instance grou
H. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.

Answer: B

Explanation:
https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer https://cloud.google.com/kubernetes-engine/docs/concepts/network-overview#ext-lb

NEW QUESTION 116
- (Topic 5)
You have deployed several instances on Compute Engine. As a security requirement, instances cannot have a public IP address. There is no VPN connection between Google
Cloud and your office, and you need to connect via SSH into a specific machine without violating the security requirements. What should you do?

A. Configure Cloud NAT on the subnet where the instance is hoste
B. Create an SSH connection to the Cloud NAT IP address to reach the instance.
C. Add all instances to an unmanaged instance grou
D. Configure TCP Proxy Load Balancing with the instance group as a backen
E. Connect to the instance using the TCP Proxy IP.
F. Configure Identity-Aware Proxy (IAP) for the instance and ensure that you have the role of IAP-secured Tunnel Use
G. Use the gcloud command line tool to ssh into the instance.
H. Create a bastion host in the network to SSH into the bastion host from your office locatio
I. From the bastion host, SSH into the desired instance.

Answer: C

Explanation:
https://cloud.google.com/iap/docs/using-tcp-forwarding#tunneling_with_ssh
Leveraging the BeyondCorp security model. "This January, we enhanced context-aware access capabilities in Cloud Identity-Aware Proxy (IAP) to help you protect SSH and RDP access to your virtual machines (VMs)—without needing to provide your VMs with public IP addresses, and without having to set up bastion hosts. "
https://cloud.google.com/blog/products/identity-security/cloud-iap-enables-context-aware- access-to-vms-via-ssh-and-rdp-without-bastion-hosts
Reference: https://cloud.google.com/solutions/connecting-securely

NEW QUESTION 119
- (Topic 5)
The database administration team has asked you to help them improve the performance of their new database server running on Google Compute Engine. The database is for importing and normalizing their performance statistics and is built with MySQL running on Debian Linux. They have an n1-standard-8 virtual machine with 80 GB of SSD persistent disk. What should they change to get better performance from this system?

A. Increase the virtual machine's memory to 64 GB.
B. Create a new virtual machine running PostgreSQL.
C. Dynamically resize the SSD persistent disk to 500 GB.
D. Migrate their performance metrics warehouse to BigQuery.
E. Modify all of their batch jobs to use bulk inserts into the database.

Answer: C

NEW QUESTION 122
- (Topic 5)
During a high traffic portion of the day, one of your relational databases crashes, but the replica is never promoted to a master. You want to avoid this in the future. What should you do?

A. Use a different database.
B. Choose larger instances for your database.
C. Create snapshots of your database more regularly.
D. Implement routinely scheduled failovers of your databases.

Answer: D

Explanation:
https://cloud.google.com/solutions/dr-scenarios-planning-guide

NEW QUESTION 126
- (Topic 5)
The operations team in your company wants to save Cloud VPN log events (or one year You need to configure the cloud infrastructure to save the logs What should you do?

A. Set up a filter in Cloud Logging and a topic in Pub/Sub to publish the logs
B. Set up a Cloud Logging Dashboard titled Cloud VPN Logs, and then add a chart that queries for the VPN metrics over a one-year time period
C. Enable the Compute Engine API and then enable logging on the firewall rules that match the traffic you want to save
D. Set up a filter in Cloud Logging and a Cloud Storage bucket as an export target for the logs you want to save

Answer: D

NEW QUESTION 127
- (Topic 5)
Your company has sensitive data in Cloud Storage buckets. Data analysts have Identity
Access Management (IAM) permissions to read the buckets. You want to prevent data analysts from retrieving the data in the buckets from outside the office network. What should you do?

A. * 1. Create a VPC Service Controls perimeter that includes the projects with the buckets.* 2. Create an access level with the CIDR of the office network.
B. * 1. Create a firewall rule for all instances in the Virtual Private Cloud (VPC) network for source range.* 2. Use the Classless Inter-domain Routing (CIDR) of the office network.
C. * 1. Create a Cloud Function to remove IAM permissions from the buckets, and another Cloud Function to add IAM permissions to the buckets.* 2. Schedule the Cloud Functions with Cloud Scheduler to add permissions at the start of business and remove permissions at the end of business.
D. * 1. Create a Cloud VPN to the office network.* 2. Configure Private Google Access for on-premises hosts.

Answer: A

Explanation:
For all Google Cloud services secured with VPC Service Controls, you can ensure that: Resources within a perimeter are accessed only from clients within authorized VPC networks using Private Google Access with either Google Cloud or on-premises. https://cloud.google.com/vpc-service-controls/docs/overview https://cloud.google.com/vpc-service-controls/docs/overview. You create a service control
across your VPC and any cloud bucket or any project resource to restrict access. Anything outside of it can't access the resources within service control perimeter

NEW QUESTION 131
- (Topic 5)
Your company's test suite is a custom C++ application that runs tests throughout each day on Linux virtual machines. The full test suite takes several hours to complete, running on a limited number of on premises servers reserved for testing. Your company wants to move the testing infrastructure to the cloud, to reduce the amount of time it takes to fully test a change to the system, while changing the tests as little as possible. Which cloud infrastructure should you recommend?

A. Google Compute Engine unmanaged instance groups and Network Load Balancer
B. Google Compute Engine managed instance groups with auto-scaling
C. Google Cloud Dataproc to run Apache Hadoop jobs to process each test
D. Google App Engine with Google Stackdriver for logging

Answer: B

Explanation:
https://cloud.google.com/compute/docs/instance-groups/
Google Compute Engine enables users to launch virtual machines (VMs) on demand. VMs can be launched from the standard images or custom images created by users.
Managed instance groups offer autoscaling capabilities that allow you to automatically add or remove instances from a managed instance group based on increases or decreases in load. Autoscaling helps your applications gracefully handle increases in traffic and reduces cost when the need for resources is lower.

NEW QUESTION 133
- (Topic 5)
Your organization has a 3-tier web application deployed in the same network on Google Cloud Platform. Each tier (web, API, and database) scales independently of the others Network traffic should flow through the web to the API tier and then on to the database tier. Traffic should not flow between the web and the database tier. How should you configure the network?

A. Add each tier to a different subnetwork.
B. Set up software based firewalls on individual VMs.
C. Add tags to each tier and set up routes to allow the desired traffic flow.
D. Add tags to each tier and set up firewall rules to allow the desired traffic flow.

Answer: D

Explanation:
https://aws.amazon.com/blogs/aws/building-three-tier-architectures-with- security-groups/
Google Cloud Platform(GCP) enforces firewall rules through rules and tags. GCP rules and tags can be defined once and used across all regions.
References: https://cloud.google.com/docs/compare/openstack/ https://aws.amazon.com/it/blogs/aws/building-three-tier-architectures-with-security-groups/

NEW QUESTION 135
- (Topic 5)
You are developing an application using different microservices that should remain internal to the cluster. You want to be able to configure each microservice with a specific number of replicas. You also want to be able to address a specific microservice from any other microservice in a uniform way, regardless of the number of replicas the microservice scales to. You need to implement this solution on Google Kubernetes Engine. What should you do?

A. Deploy each microservice as a Deploymen
B. Expose the Deployment in the cluster using a Service, and use the Service DNS name to address it from other microservices within the cluster.

C. Deploy each microservice as a Deploymen
D. Expose the Deployment in the cluster using an Ingress, and use the Ingress IP address to address the Deployment from other microservices within the cluster.
E. Deploy each microservice as a Po
F. Expose the Pod in the cluster using a Service, anduse the Service DNS name to address the microservice from other microservices within the cluster.
G. Deploy each microservice as a Po
H. Expose the Pod in the cluster using an Ingress, and use the Ingress IP address name to address the Pod from other microservices within the cluster.

Answer: A

Explanation:
https://kubernetes.io/docs/concepts/services-networking/ingress/

NEW QUESTION 136
- (Topic 5)
You are designing a mobile chat application. You want to ensure people cannot spoof chat messages, by providing a message were sent by a specific user. What should you do

A. Tag messages client side with the originating user identifier and the destination user.
B. Encrypt the message client side using block-based encryption with a shared key.
C. Use public key infrastructure (PKI) to encrypt the message client side using the originating user's privatekey.
D. Use a trusted certificate authority to enable SSL connectivity between the client application and the server.

Answer: C

NEW QUESTION 141
- (Topic 5)
Your company has just acquired another company, and you have been asked to integrate their existing Google Cloud environment into your company's data center. Upon investigation, you discover that some of the RFC 1918 IP ranges being used in the new company's Virtual Private Cloud (VPC) overlap with your data center IP space. What should you do to enable connectivity and make sure that there are no routing conflicts when connectivity is established?

A. Create a Cloud VPN connection from the new VPC to the data center, create a Cloud Router, and apply new IP addresses so there is no overlapping IP space.
B. Create a Cloud VPN connection from the new VPC to the data center, and create a Cloud NAT instance to perform NAT on the overlapping IP space.
C. Create a Cloud VPN connection from the new VPC to the data center, create a Cloud Router, and apply a custom route advertisement to block the overlapping IP space.
D. Create a Cloud VPN connection from the new VPC to the data center, and apply a firewall rule that blocks the overlapping IP space.

Answer: A

Explanation:
To connect two networks together we need (1) either VPN or interconnect and (2) peering. When there is peering, you cannot have conflicting IP addresses. You can use either Cloud VPN or Cloud Interconnect to securely connect your on-premises network to your VPC network. (https://cloud.google.com/vpc/docs/vpc-peering#transit-network) At the time of peering, Google Cloud checks to see if there are any subnet IP ranges that overlap subnet IP ranges in the other network. If there is any overlap, peering is not established. (https://cloud.google.com/vpc/docs/vpc-peering#considerations) NAT is used to translate private to public IP and vice versa, however because we are connecting 2 networks together, they become private IPs. So it is not applicable.

NEW QUESTION 142
- (Topic 5)
You want to create a private connection between your instances on Compute Engine and your on-premises data center. You require a connection of at least 20 Gbps. You want to follow Google-recommended practices.
How should you set up the connection?

A. Create a VPC and connect it to your on-premises data center using Dedicated Interconnect.
B. Create a VPC and connect it to your on-premises data center using a single Cloud VPN.
C. Create a Cloud Content Delivery Network (Cloud CDN) and connect it to your on- premises data centerusing Dedicated Interconnect.
D. Create a Cloud Content Delivery Network (Cloud CDN) and connect it to your on- premises datacenterusing a single Cloud VPN.

Answer: A

Explanation:
Reference: https://cloud.google.com/compute/docs/instances/connecting-advanced

NEW QUESTION 144
- (Topic 5)
You want to make a copy of a production Linux virtual machine in the US-Central region. You want to manage and replace the copy easily if there are changes on the production
virtual machine. You will deploy the copy as a new instances in a different project in the US-East region. What steps must you take?

A. Use the Linux dd and netcat command to copy and stream the root disk contents to a new virtual machine instance in the US-East region.
B. Create a snapshot of the root disk and select the snapshot as the root disk when you create a new virtual machine instance in the US-East region.
C. Create an image file from the root disk with Linux dd command, create a new disk from the image file, and use it to create a new virtual machine instance in the US-East region
D. Create a snapshot of the root disk, create an image file in Google Cloud Storage from the snapshot, and create a new virtual machine instance in the US-East region using the image file for the root disk.

Answer: D

Explanation:
https://stackoverflow.com/questions/36441423/migrate-google-compute-engine-instance-to-a-different-region

NEW QUESTION 146
- (Topic 6)
For this question, refer to the Dress4Win case study. Considering the given business requirements, how would you automate the deployment of web and transactional data layers?

A. Deploy Nginx and Tomcat using Cloud Deployment Manager to Compute Engin
B. Deploy a Cloud SQL server to replace MySQ
C. Deploy Jenkins using Cloud Deployment Manager.
D. Deploy Nginx and Tomcat using Cloud Launche
E. Deploy a MySQL server using Cloud Launche
F. Deploy Jenkins to Compute Engine using Cloud Deployment Manager scripts.
G. Migrate Nginx and Tomcat to App Engin
H. Deploy a Cloud Datastore server to replace the MySQL server in a high-availability configuratio
I. Deploy Jenkins to Compute Engine using Cloud Launcher.
J. Migrate Nginx and Tomcat to App Engin
K. Deploy a MySQL server using Cloud Launche
L. Deploy Jenkins to Compute Engine using Cloud Launcher.

Answer: A

NEW QUESTION 147
- (Topic 6)
For this question, refer to the Dress4Win case study. Dress4Win is expected to grow to 10 times its size in 1 year with a corresponding growth in data and traffic that mirrors the existing patterns of usage. The CIO has set the target of migrating production infrastructure to the cloud within the next 6 months. How will you configure the solution to scale for this growth without making major application changes and still maximize the ROI?

A. Migrate the web application layer to App Engine, and MySQL to Cloud Datastore, and NAS to Cloud Storag
B. Deploy RabbitMQ, and deploy Hadoop servers using Deployment Manager.
C. Migrate RabbitMQ to Cloud Pub/Sub, Hadoop to BigQuery, and NAS to Compute Engine with Persistent Disk storag
D. Deploy Tomcat, and deploy Nginx using Deployment Manager.
E. Implement managed instance groups for Tomcat and Ngin
F. Migrate MySQL to Cloud SQL, RabbitMQ to Cloud Pub/Sub, Hadoop to Cloud Dataproc, and NAS to Compute Engine with Persistent Disk storage.
G. Implement managed instance groups for the Tomcat and Ngin
H. Migrate MySQL to Cloud SQL, RabbitMQ to Cloud Pub/Sub, Hadoop to Cloud Dataproc, and NAS to Cloud Storage.

Answer: D

NEW QUESTION 152
- (Topic 7)
For this question, refer to the TerramEarth case study. Considering the technical requirements, how should you reduce the unplanned vehicle downtime in GCP?

A. Use BigQuery as the data warehous
B. Connect all vehicles to the network and stream data into BigQuery using Cloud Pub/Sub and Cloud Dataflo
C. Use Google Data Studio for analysis and reporting.
D. Use BigQuery as the data warehous
E. Connect all vehicles to the network and upload gzip files to a Multi-Regional Cloud Storage bucket using gclou
F. Use Google Data Studio for analysis and reporting.
G. Use Cloud Dataproc Hive as the data warehous
H. Upload gzip files to a MultiRegional Cloud Storagebucke
I. Upload this data into BigQuery using gclou
J. Use Google data Studio for analysis and reporting.
K. Use Cloud Dataproc Hive as the data warehous
L. Directly stream data into prtitioned Hive table
M. Use Pig scripts to analyze data.

Answer: A

NEW QUESTION 157
- (Topic 7)
TerramEarth has about 1 petabyte (PB) of vehicle testing data in a private data center. You want to move the data to Cloud Storage for your machine learning team. Currently, a 1- Gbps interconnect link is available for you. The machine learning team wants to start using
the data in a month. What should you do?

A. Request Transfer Appliances from Google Cloud, export the data to appliances, and return the appliances to Google Cloud.
B. Configure the Storage Transfer service from Google Cloud to send the data from your data center to Cloud Storage
C. Make sure there are no other users consuming the 1 Gbps link, and use multi-thread transfer to upload the data to Cloud Storage.
D. Export files to an encrypted USB device, send the device to Google Cloud, and request an import of the data to Cloud Storage

Answer: A

NEW QUESTION 160
- (Topic 8)
For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to migrate from their current analytics and statistics reporting model to one that meets their technical requirements on Google Cloud Platform.
Which two steps should be part of their migration plan? (Choose two.)

A. Evaluate the impact of migrating their current batch ETL code to Cloud Dataflow.
B. Write a schema migration plan to denormalize data for better performance in BigQuery.
C. Draw an architecture diagram that shows how to move from a single MySQL database to a MySQL cluster.

D. Load 10 TB of analytics data from a previous game into a Cloud SQL instance, and run test queries against the full dataset to confirm that they complete successfully.
E. Integrate Cloud Armor to defend against possible SQL injection attacks in analytics files uploaded to Cloud Storage.

Answer: AB

Explanation:
https://cloud.google.com/bigquery/docs/loading-data#loading_denormalized_nested_and_repeated_data

NEW QUESTION 162
- (Topic 8)
Mountkirk Games wants you to secure the connectivity from the new gaming application platform to Google Cloud. You want to streamline the process and follow Google-recommended practices. What should you do?

A. Configure Workload Identity and service accounts to be used by the application platform.
B. Use Kubernetes Secrets, which are obfuscated by defaul
C. Configure these Secrets to be used by theapplication platform.
D. Configure Kubernetes Secrets to store the secret, enable Application-Layer Secrets Encryption, and useCloud Key Management Service (Cloud KMS) to manage the encryption key
E. Configure these Secrets tobe used by the application platform.
F. Configure HashiCorp Vault on Compute Engine, and use customer managed encryption keys and CloudKey Management Service (Cloud KMS) to manage the encryption key
G. Configure these Secrets to be usedby the application platform.

Answer: A

NEW QUESTION 163
- (Topic 8)
For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the compute workloads for your company, Mountkirk Games. Considering the Mountkirk Games business and technical requirements, what should you do?

A. Create network load balancer
B. Use preemptible Compute Engine instances.
C. Create network load balancer
D. Use non-preemptible Compute Engine instances.
E. Create a global load balancer with managed instance groups and autoscaling policie
F. Use preemptible Compute Engine instances.
G. Create a global load balancer with managed instance groups and autoscaling policie
H. Use non-preemptible Compute Engine instances.

Answer: D

NEW QUESTION 165
- (Topic 9)
For this question, refer to the Helicopter Racing League (HRL) case study. The HRL development team releases a new version of their predictive capability application every Tuesday evening at 3

A. a.
B. UTC to a repositor
C. The security team at HRL has developed an in-house penetration test Cloud Function called Airwolf.The security team wants to run Airwolf against the predictive capability application as soon as it is releasedevery Tuesda
D. You need to set up Airwolf to run at the recurring weekly cadenc
E. What should you do?
F. Set up Cloud Tasks and a Cloud Storage bucket that triggers a Cloud Function.
G. Set up a Cloud Logging sink and a Cloud Storage bucket that triggers a Cloud Function.
H. Configure the deployment job to notify a Pub/Sub queue that triggers a Cloud Function.
I. Set up Identity and Access Management (IAM) and Confidential Computing to trigger a Cloud Function.

Answer: A

NEW QUESTION 166
- (Topic 9)
For this question, refer to the Helicopter Racing League (HRL) case study. HRL wants better prediction accuracy from their ML prediction models. They want you to use Google's AI Platform so HRL can understand and interpret the predictions. What should you do?

A. Use Explainable AI.
B. Use Vision AI.
C. Use Google Cloud's operations suite.
D. Use Jupyter Notebooks.

Answer: A

Explanation:
Reference: https://cloud.google.com/ai-platform/prediction/docs/ai-explanations/preparing- metadata

NEW QUESTION 171
- (Topic 10)
For this question, refer to the EHR Healthcare case study. In the past, configuration errors put public IP addresses on backend servers that should not have been

accessible from the Internet. You need to ensure that no one can put external IP addresses on backend Compute Engine instances and that external IP addresses can only be configured on frontend Compute Engine instances. What should you do?

A. Create an Organizational Policy with a constraint to allow external IP addresses only on the frontend Compute Engine instances.
B. Revoke the compute.networkAdmin role from all users in the project with front end instances.
C. Create an Identity and Access Management (IAM) policy that maps the IT staff to the compute.networkAdmin role for the organization.
D. Create a custom Identity and Access Management (IAM) role named GCE_FRONTEND with the compute.addresses.create permission.

Answer: A

Explanation:
https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#disableexternalip

NEW QUESTION 174
- (Topic 10)
For this question, refer to the EHR Healthcare case study. EHR has single Dedicated Interconnect connection between their primary data center and Googles network. This connection satisfies EHR's network and security policies:
• On-premises servers without public IP addresses need to connect to cloud resources without public IP addresses
• Traffic flows from production network mgmt. servers to Compute Engine virtual machines should never traverse the public internet.
You need to upgrade the EHR connection to comply with their requirements. The new connection design must support business critical needs and meet the same network and security policy requirements. What should you do?

A. Add a new Dedicated Interconnect connection
B. Upgrade the bandwidth on the Dedicated Interconnect connection to 100 G
C. Add three new Cloud VPN connections
D. Add a new Carrier Peering connection

Answer: A

Explanation:
The case does not call out the throughput being an issue. However, to achieve 99.99%, you need to have 4 connections as per Google recommendations. However, in the options only A has the option to add an additional Interconnect connection. https://cloud.google.com/network-connectivity/docs/interconnect/concepts/dedicated- overview#availability

NEW QUESTION 175
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

Professional-Cloud-Architect Practice Exam Features:

* Professional-Cloud-Architect Questions and Answers Updated Frequently

* Professional-Cloud-Architect Practice Questions Verified by Expert Senior Certified Staff

* Professional-Cloud-Architect Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* Professional-Cloud-Architect Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year