Professional-Cloud-Architect Dumps

Google Certified Professional - Cloud Architect (GCP)

https://www.certleader.com/Professional-Cloud-Architect-dumps.html

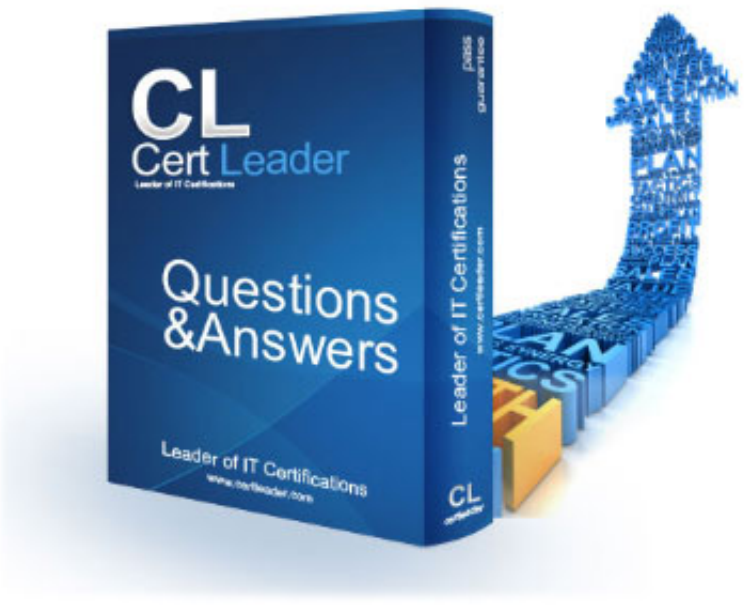

NEW QUESTION 1
- (Topic 1)
For this question, refer to the Mountkirk Games case study.
Mountkirk Games wants to set up a continuous delivery pipeline. Their architecture includes many small services that they want to be able to update and roll back quickly. Mountkirk Games has the following requirements:
• Services are deployed redundantly across multiple regions in the US and Europe.
• Only frontend services are exposed on the public internet.
• They can provide a single frontend IP for their fleet of services.
• Deployment artifacts are immutable. Which set of products should they use?

A. Google Cloud Storage, Google Cloud Dataflow, Google Compute Engine
B. Google Cloud Storage, Google App Engine, Google Network Load Balancer
C. Google Kubernetes Registry, Google Container Engine, Google HTTP(S) Load Balancer
D. Google Cloud Functions, Google Cloud Pub/Sub, Google Cloud Deployment Manager

Answer: C

NEW QUESTION 2
- (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 3
- (Topic 1)
For this question, refer to the Mountkirk Games case study.
Mountkirk Games wants you to design their new testing strategy. How should the test coverage differ from their existing backends on the other platforms?

A. Tests should scale well beyond the prior approaches.
B. Unit tests are no longer required, only end-to-end tests.
C. Tests should be applied after the release is in the production environment.
D. Tests should include directly testing the Google Cloud Platform (GCP) infrastructure.

Answer: A

Explanation:
From Scenario:
A few of their games were more popular than expected, and they had problems scaling their application servers, MySQL databases, and analytics tools.
Requirements for Game Analytics Platform include: Dynamically scale up or down based on game activity

NEW QUESTION 4
- (Topic 2)
For this question, refer to the TerramEarth case study.
TerramEarth's CTO wants to use the raw data from connected vehicles to help identify approximately when a vehicle in the development team to focus their failure.
You want to allow analysts to centrally query the vehicle data. Which architecture should you recommend?
A)

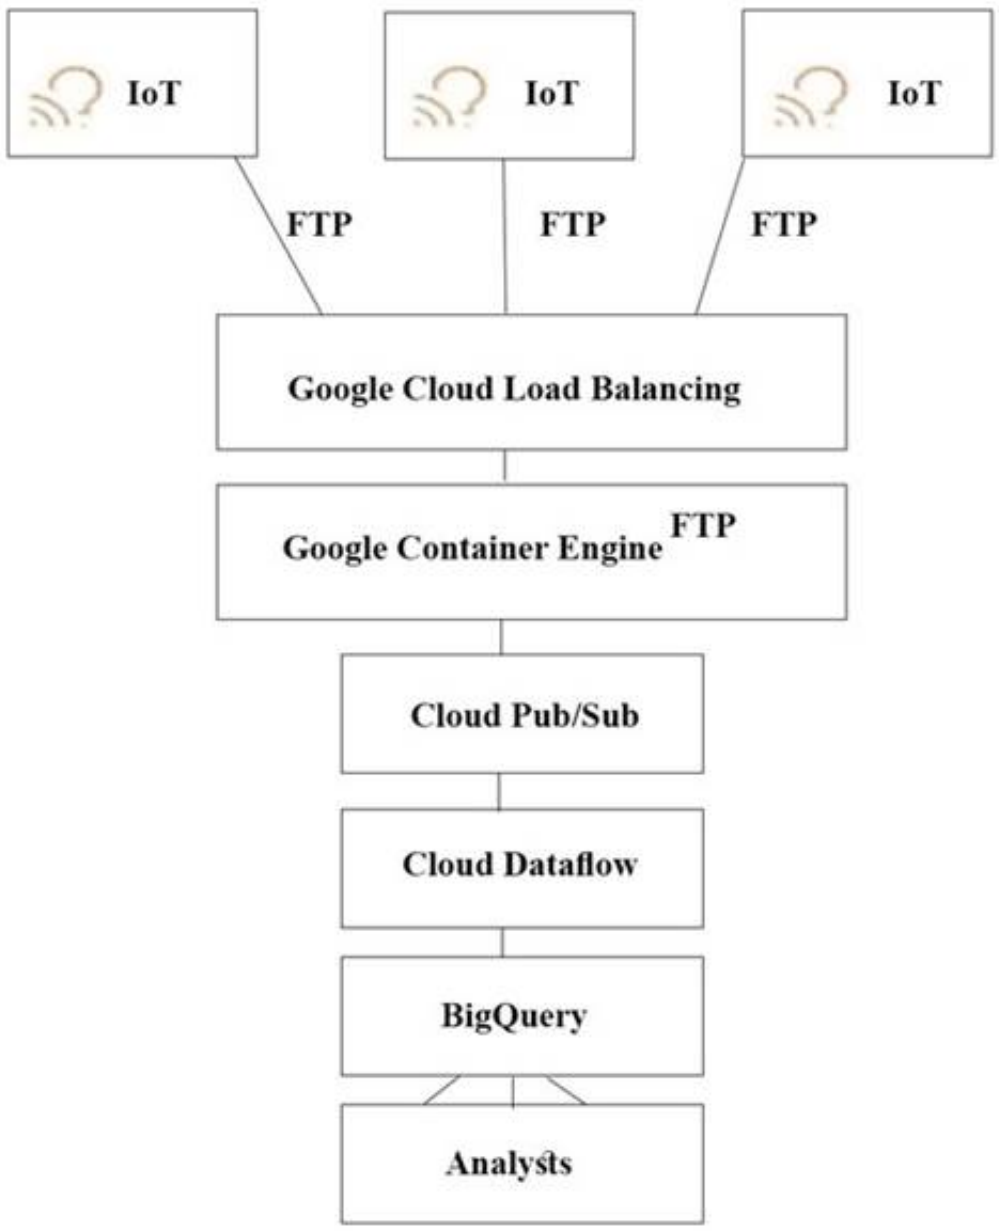
IoT
IoT
IoT
FTP
FTP
FTP
Google Cloud Load Balancing
Google Container Engine
FTP
Cloud Pub/Sub
Cloud Dataflow
BigQuery
Analysts

B)

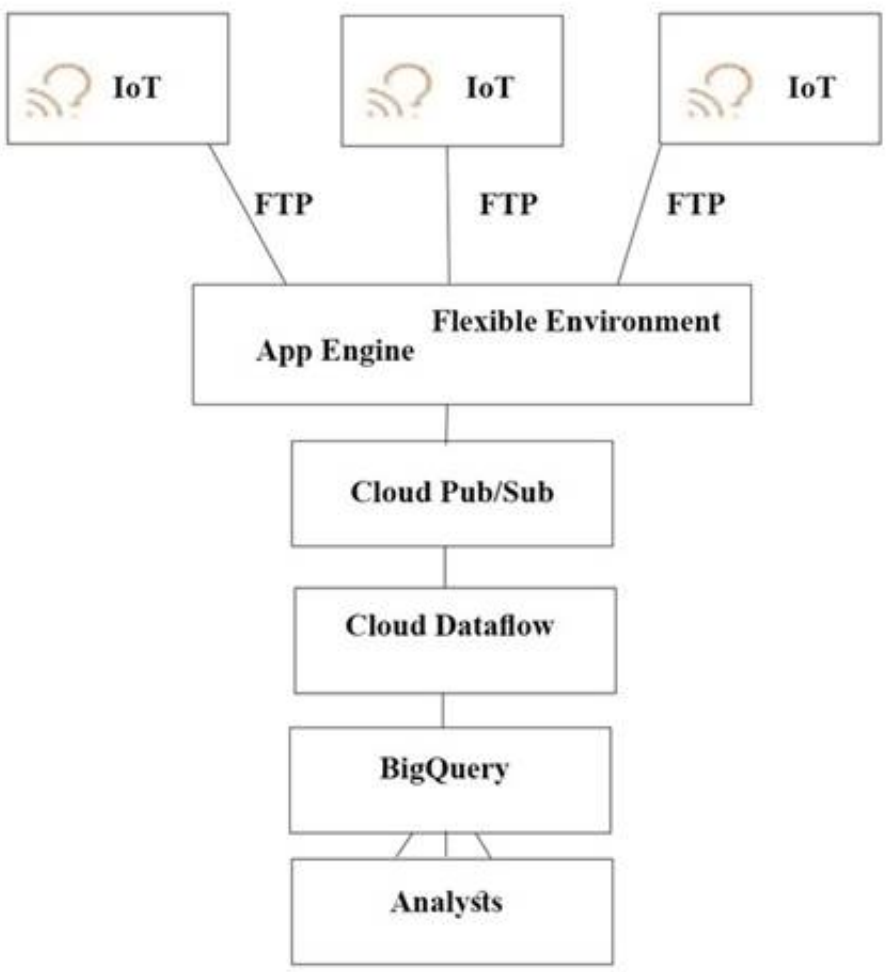
IoT
IoT
IoT
FTP
FTP
FTP
App Engine
Flexible Environment
Cloud Pub/Sub
Cloud Dataflow
BigQuery
Analysts

C)

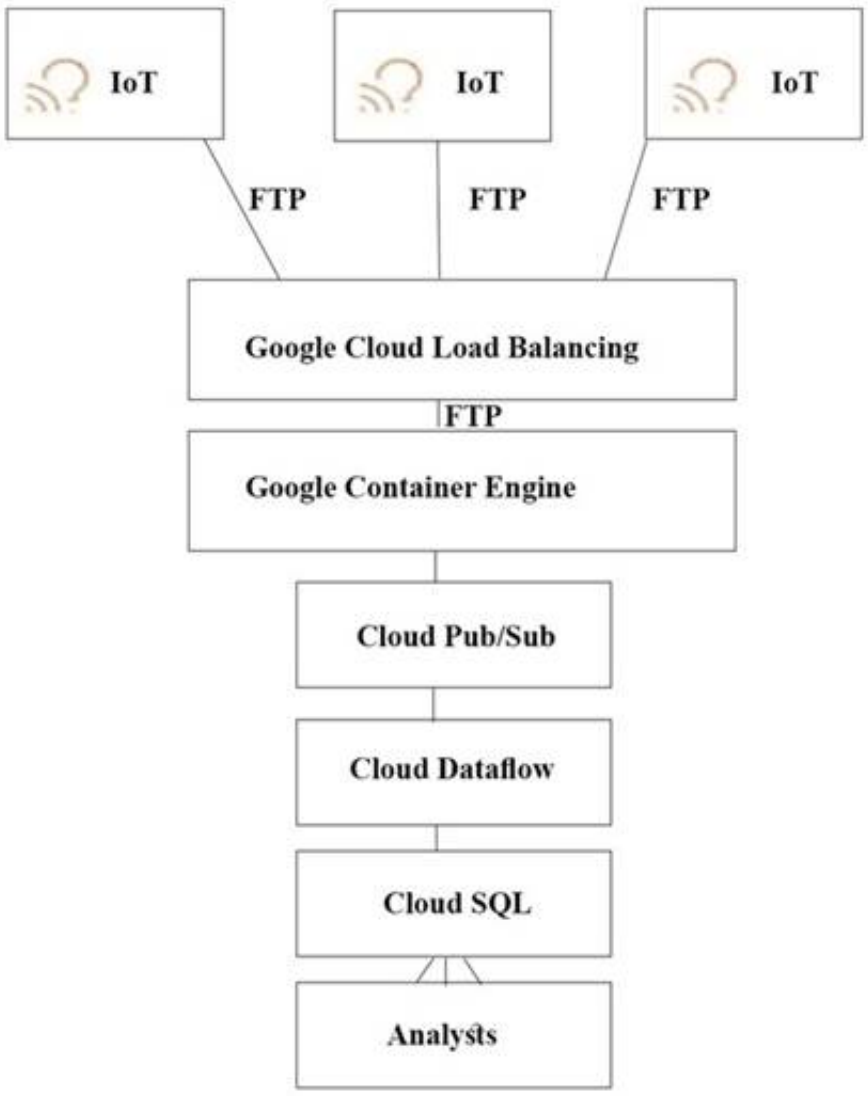

D)

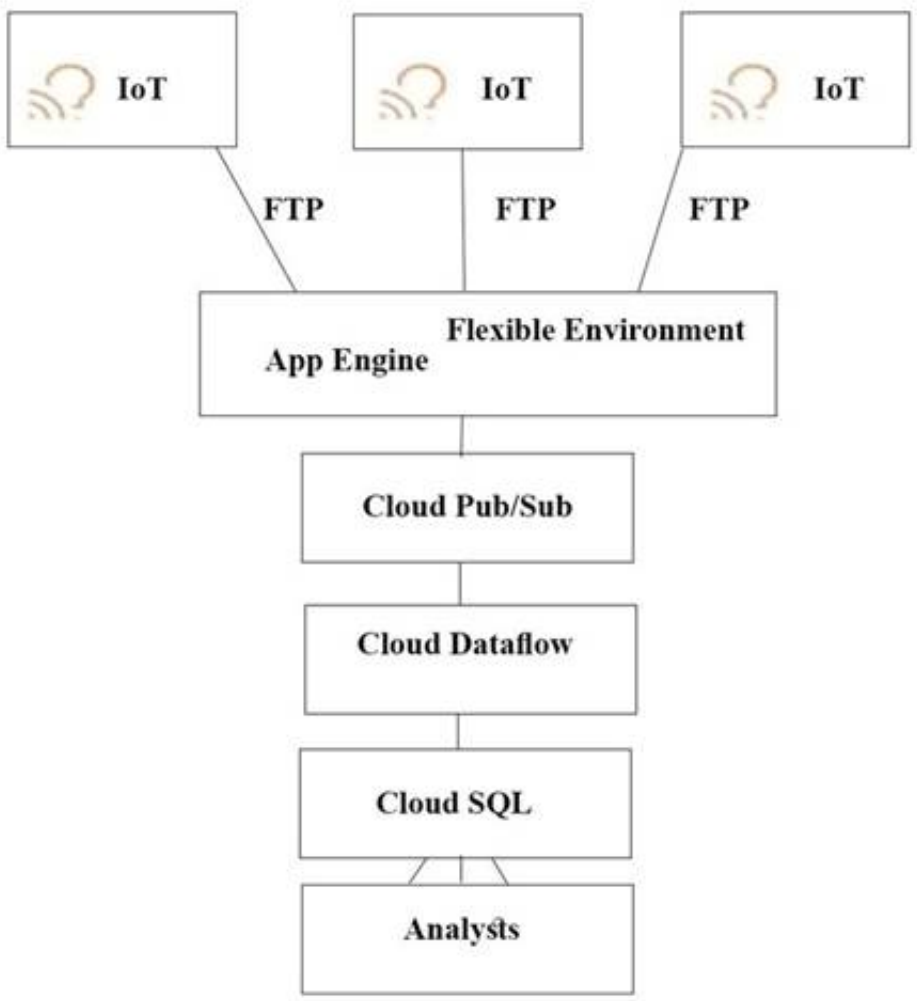

A. Option A
B. Option B
C. Option C
D. Option D

Answer: A

Explanation:
https://cloud.google.com/solutions/iot/ https://cloud.google.com/solutions/designing-connected-vehicle-platform https://cloud.google.com/solutions/designing-connected-vehicle-platform#data_ingestion
http://www.eweek.com/big-data-and-analytics/google-touts-value-of-cloud-iot-core-for- analyzing-connected-car-data
https://cloud.google.com/solutions/iot/
The push endpoint can be a load balancer.
A container cluster can be used. Cloud Pub/Sub for Stream Analytics

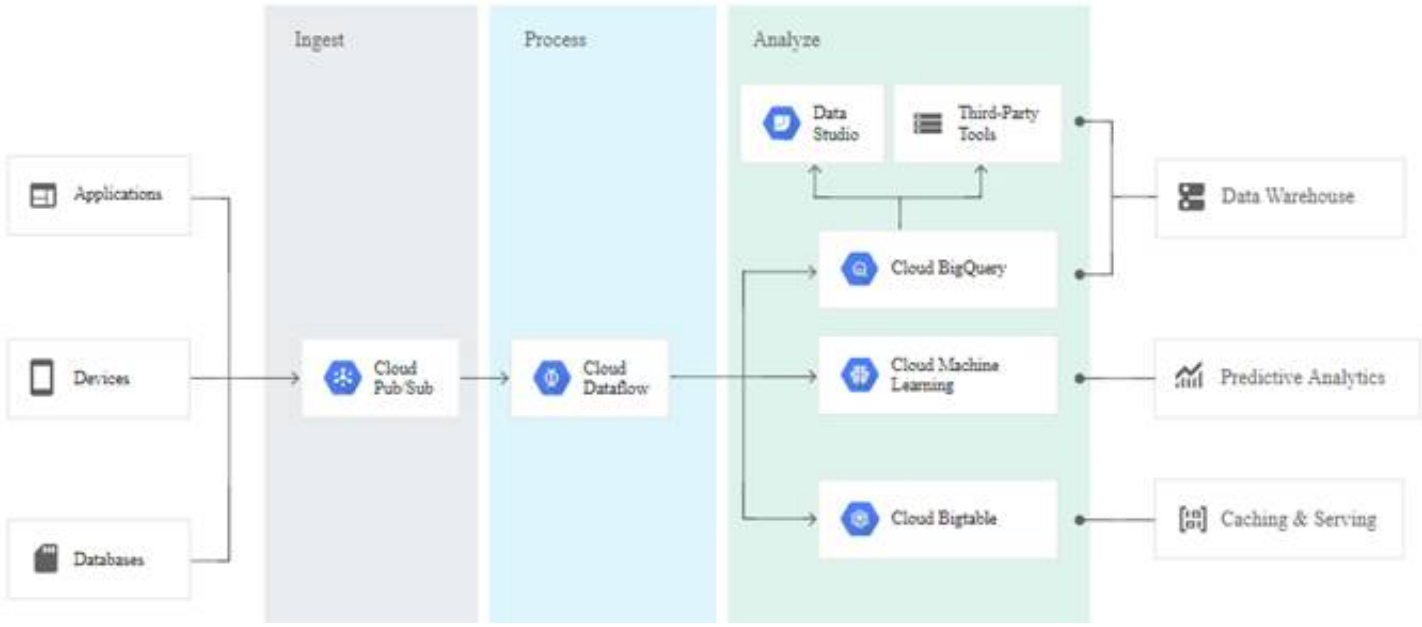

References: https://cloud.google.com/pubsub/ https://cloud.google.com/solutions/iot/ https://cloud.google.com/solutions/designing-connected-vehicle-platform https://cloud.google.com/solutions/designing-connected-vehicle-platform#data_ingestion http://www.eweek.com/big-data-and-analytics/google-touts-value-of-cloud-iot-core-for-analyzing-connected-car-data
https://cloud.google.com/solutions/iot/

NEW QUESTION 5
- (Topic 2)
For this question, refer to the TerramEarth case study.
TerramEarth has equipped unconnected trucks with servers and sensors to collet telemetry data. Next year they want to use the data to train machine learning models. They want to store this data in the cloud while reducing costs. What should they do?

A. Have the vehicle' computer compress the data in hourly snapshots, and store it in a Google Cloud storage (GCS) Nearline bucket.
B. Push the telemetry data in Real-time to a streaming dataflow job that compresses the data, and store it in Google BigQuery.
C. Push the telemetry data in real-time to a streaming dataflow job that compresses the data, and store it in Cloud Bigtable.
D. Have the vehicle's computer compress the data in hourly snapshots, a Store it in a GCS Coldline bucket.

Answer: D

Explanation:
Coldline Storage is the best choice for data that you plan to access at most once a year, due to its slightly lower availability, 90-day minimum storage duration, costs for data access, and higher per-operation costs. For example:
Cold Data Storage - Infrequently accessed data, such as data stored for legal or regulatory reasons, can be stored at low cost as Coldline Storage, and be available when you need it. Disaster recovery - In the event of a disaster recovery event, recovery time is key. Cloud Storage provides low latency access to data stored as Coldline Storage.
References: https://cloud.google.com/storage/docs/storage-classes

NEW QUESTION 6
- (Topic 2)
For this question, refer to the TerramEarth case study.
TerramEarth plans to connect all 20 million vehicles in the field to the cloud. This increases the volume to 20 million 600 byte records a second for 40 TB an hour. How should you design the data ingestion?

A. Vehicles write data directly to GCS.
B. Vehicles write data directly to Google Cloud Pub/Sub.
C. Vehicles stream data directly to Google BigQuery.
D. Vehicles continue to write data using the existing system (FTP).

Answer: B

Explanation:
https://cloud.google.com/solutions/data-lifecycle-cloud-platform https://cloud.google.com/solutions/designing-connected-vehicle-platform

NEW QUESTION 7
- (Topic 2)
For this question refer to the TerramEarth case study.
Which of TerramEarth's legacy enterprise processes will experience significant change as a result of increased Google Cloud Platform adoption.

A. Opex/capex allocation, LAN changes, capacity planning
B. Capacity planning, TCO calculations, opex/capex allocation
C. Capacity planning, utilization measurement, data center expansion
D. Data Center expansion, TCO calculations, utilization measurement

Answer: B

Explanation:
Capacity planning, TCO calculations, opex/capex allocation From the case study, it can conclude that Management (CXO) all concern rapid provision of resources (infrastructure)
for growing as well as cost management, such as Cost optimization in Infrastructure, trade up front capital expenditures (Capex) for ongoing operating expenditures (Opex), and Total cost of ownership (TCO)

NEW QUESTION 8
- (Topic 2)
Your agricultural division is experimenting with fully autonomous vehicles.
You want your architecture to promote strong security during vehicle operation. Which two architecture should you consider?
Choose 2 answers:

A. Treat every micro service call between modules on the vehicle as untrusted.
B. Require IPv6 for connectivity to ensure a secure address space.
C. Use a trusted platform module (TPM) and verify firmware and binaries on boot.
D. Use a functional programming language to isolate code execution cycles.
E. Use multiple connectivity subsystems for redundancy.
F. Enclose the vehicle's drive electronics in a Faraday cage to isolate chips.

Answer: AC

NEW QUESTION 9
- (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 10
- (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 10
- (Topic 4)
For this question, refer to the Dress4Win case study.
The Dress4Win security team has disabled external SSH access into production virtual machines (VMs) on Google Cloud Platform (GCP). The operations team needs to remotely manage the VMs, build and push Docker containers, and manage Google Cloud Storage objects. What can they do?

A. Grant the operations engineers access to use Google Cloud Shell.
B. Configure a VPN connection to GCP to allow SSH access to the cloud VMs.
C. Develop a new access request process that grants temporary SSH access to cloud VMs when an operations engineer needs to perform a task.
D. Have the development team build an API service that allows the operations team to execute specific remote procedure calls to accomplish their tasks.

Answer: A

NEW QUESTION 13
- (Topic 4)
For this question, refer to the Dress4Win case study.
Dress4Win has end-to-end tests covering 100% of their endpoints. They want to ensure that the move to the cloud does not introduce any new bugs. Which additional testing methods should the developers employ to prevent an outage?

A. They should enable Google Stackdriver Debugger on the application code to show errors in the code.
B. They should add additional unit tests and production scale load tests on their cloud staging environment.
C. They should run the end-to-end tests in the cloud staging environment to determine if the code is working as intended.
D. They should add canary tests so developers can measure how much of an impact the new release causes to latency.

Answer: B

NEW QUESTION 15
- (Topic 4)
For this question, refer to the Dress4Win case study.
As part of Dress4Win's plans to migrate to the cloud, they want to be able to set up a managed logging and monitoring system so they can handle spikes in their traffic load.
They want to ensure that:
• The infrastructure can be notified when it needs to scale up and down to handle the ebb and flow of usage throughout the day
• Their administrators are notified automatically when their application reports errors.
• They can filter their aggregated logs down in order to debug one piece of the application across many hosts
Which Google StackDriver features should they use?

A. Logging, Alerts, Insights, Debug
B. Monitoring, Trace, Debug, Logging
C. Monitoring, Logging, Alerts, Error Reporting
D. Monitoring, Logging, Debug, Error Report

Answer: D

NEW QUESTION 17
- (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 18
- (Topic 4)
For this question, refer to the Dress4Win case study.
Dress4Win has asked you for advice on how to migrate their on-premises MySQL deployment to the cloud. They want to minimize downtime and performance impact to their on-premises solution during the migration. Which approach should you recommend?

A. Create a dump of the on-premises MySQL master server, and then shut it down, upload it to the cloud environment, and load into a new MySQL cluster.
B. Setup a MySQL replica server/slave in the cloud environment, and configure it for asynchronous replication from the MySQL master server on-premises until cutover.
C. Create a new MySQL cluster in the cloud, configure applications to begin writing to both on-premises and cloud MySQL masters, and destroy the original cluster at cutover.
D. Create a dump of the MySQL replica server into the cloud environment, load it into: Google Cloud Datastore, and configure applications to read/write to Cloud Datastore at cutover.

Answer: B

NEW QUESTION 22
- (Topic 5)
Your company captures all web traffic data in Google Analytics 260 and stores it in BigQuery. Each country has its own dataset. Each dataset has multiple tables.
You want analysts from each country
to be able to see and query only the data for their respective countries. How should you configure the access rights?

A. Create a group per countr
B. Add analysts to their respective country-group
C. Create asingle group‘all_analysts’, and add all country-groups as member
D. Grant the ‘all-analysis’ group the IAM role ofBigQuery jobUse
E. Share the appropriate dataset with view access with each respective analyst country-group.
F. Create a group per countr
G. Add analysts to their respective country-group

H. Create a single group‘all_analysts’, and add all country-groups as member
I. Grant the ‘all-analysis’ group the IAM role ofBigQuery jobUse
J. Share the appropriate tables with view access with each respective analyst countrygroup.
K. Create a group per countr
L. Add analysts to their respective country-group
M. Create a single group‘all_analysts’, and add all country-groups as member
N. Grant the ‘all-analysis’ group the IAM role ofBigQuery dataViewe
O. Share the appropriate dataset with view access with each respective analystcountry-group.
P. Create a group per countr
Q. Add analysts to their respective country-group
R. Create a single group‘all_analysts’, and add all country-groups as member
S. Grant the ‘all-analysis’ group the IAM role ofBigQuery dataViewe
T. Share the appropriate table with view access with each respective analyst countrygroup.

Answer: A

NEW QUESTION 26
- (Topic 5)
Your company is running its application workloads on Compute Engine. The applications have been deployed in production, acceptance, and development environments. The production environment is business-critical and is used 24/7, while the acceptance and development environments are only critical during office hours. Your CFO has asked you to optimize these environments to achieve cost savings during idle times. What should you do?

A. Create a shell script that uses the gcloud command to change the machine type of the development and acceptance instances to a smaller machine type outside of office hour
B. Schedule the shell script on one of the production instances to automate the task.
C. Use Cloud Scheduler to trigger a Cloud Function that will stop the development and acceptance environments after office hours and start them just before office hours.
D. Deploy the development and acceptance applications on a managed instance group and enable autoscaling.
E. Use regular Compute Engine instances for the production environment, and usepreemptible VMs for the acceptance and development environments.

Answer: B

Explanation:
Reference: https://cloud.google.com/blog/products/it-ops/best-practices-for-optimizing- your-cloud-costs

NEW QUESTION 29
- (Topic 5)
Your company has an application running on Compute Engine mat allows users to play their favorite music. There are a fixed number of instances Files are stored in Cloud Storage and data is streamed directly to users. Users are reporting that they sometimes need to attempt to play popular songs multiple times before they are successful. You need to improve the performance of the application. What should you do?
A.
* 1. Copy popular songs into CloudSQL as a blob
* 2. Update application code to retrieve data from CloudSQL when Cloud Storage is overloaded
B.
* 1. Create a managed instance group with Compute Engine instances
* 2. Create a global toad balancer and configure ii with two backbends
* Managed instance group
* Cloud Storage bucket
* 3. Enable Cloud CDN on the bucket backend
C.
* 1. Mount the Cloud Storage bucket using gcsfuse on all backend Compute Engine instances
* 2. Serve muse files directly from the backend Compute Engine instance
D.
* 1. Create a Cloud Filestore NFS volume and attach it to the backend Compute Engine instances
* 2. Download popular songs in Cloud Filestore
* 3. Serve music Wes directly from the backend Compute Engine instance

A.

Answer: B

NEW QUESTION 34
- (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 37
- (Topic 5)
Your company has a networking team and a development team. The development team runs applications on Compute Engine instances that contain sensitive data. The development team requires administrative permissions for Compute Engine. Your company requires all network resources to be managed by the networking team. The development team does not want the networking team to have access to the sensitive data on the instances. What should you do?

A. * 1. Create a project with a standalone VPC and assign the Network Admin role to the networking team.* 2. Create a second project with a standalone VPC and assign the Compute Admin role to the development team.* 3. Use Cloud VPN to join the two VPCs.
B. * 1. Create a project with a standalone Virtual Private Cloud (VPC), assign the Network Admin role to the networking team, and assign the Compute Admin role to the development team.
C. * 1. Create a project with a Shared VPC and assign the Network Admin role to the networking team.* 2. Create a second project without a VPC, configure it as a Shared VPC service project,and assign the Compute Admin role to the development team.
D. * 1. Create a project with a standalone VPC and assign the Network Admin role to the networking team.* 2. Create a second project with a standalone VPC and assign the Compute Admin role to the development team.* 3. Use VPC Peering to join the two VPCs.

Answer: C

Explanation:
In this scenario, a large organization has a central team that manages security and networking controls for the entire organization. Developers do not have permissions to make changes to any network or security settings defined by the security and networking team but they are granted permission to create resources such as virtual machines in shared subnets. To facilitate this the organization makes use of a shared VPC (Virtual Private Cloud). A shared VPC allows creation of a VPC network of RFC 1918 IP spaces that associated projects (service projects) can then use. Developers using the associated projects can create VM instances in the shared VPC network spaces. The organization's network and security admins can create subnets, VPNs, and firewall rules usable by all the projects in the VPC network. https://cloud.google.com/iam/docs/job-functions/networking#single_team_manages_security_network_for_organization
Reference: https://cloud.google.com/vpc/docs/shared-vpc

NEW QUESTION 40
- (Topic 5)
You want to store critical business information in Cloud Storage buckets. The information is regularly changed but previous versions need to be referenced on a regular basis. You want to ensure that there is a record of all changes to any information in these buckets. You want to ensure that accidental edits or deletions can be easily roiled back. Which feature should you enable?

A. Bucket Lock
B. Object Versioning
C. Object change notification
D. Object Lifecycle Management

Answer: B

NEW QUESTION 44
- (Topic 5)
Your company recently acquired a company that has infrastructure in Google Cloud. Each company has its own Google Cloud organization Each company is using a Shared Virtual Private Cloud (VPC) to provide network connectivity tor its applications Some of the subnets used by both companies overlap In order for both businesses to integrate, the applications need to have private network connectivity. These applications are not on overlapping subnets. You want to provide connectivity with minimal re-engineering. What should you do?

A. Set up VPC peering and peer each Shared VPC together
B. Configure SSH port forwarding on each application to provide connectivity between applications i the different Shared VPCs
C. Migrate the protects from the acquired company into your company's Google Cloud organization Re launch the instances in your companies Shared VPC
D. Set up a Cloud VPN gateway in each Shared VPC and peer Cloud VPNs

Answer: B

NEW QUESTION 46
- (Topic 5)
Your company has a Google Workspace account and Google Cloud Organization Some developers in the company have created Google Cloud projects outside of the Google Cloud Organization
You want to create an Organization structure that allows developers to create projects, but prevents them from modifying production projects You want to manage policies for all projects centrally and be able to set more restrictive policies for production projects
You want to minimize disruption to users and developers when business needs change in the future You want to follow Google-recommended practices How should you design the Organization structure?

A. * 1 Create a second Google Workspace account and Organization* 2 Grant all developers the Project Creator IAM role on the new Organization * 3 Move the developer projects into the new Organization* 4 Set the policies for all projects on both Organizations.* 5 Additionally set the production policies on the original Organization
B. * 1 Create a folder under the Organization resource named "Production '* 2 Grant all developers the Project Creator IAM role on the Organization * 3. Move the developer projects into the Organization* 4 Set the policies for all projects on the Organization* 5 Additionally set the production policies on the 'Production" folder
C. * 1 Create folders under the Organization resource named "Development" and Production'* 2 Grant all developers the Project Creator IAM role on the ""Development1 folder * 3. Move the developer projects into the "Development" folder* 4 Set the policies for all projects on the Organization* 5 Additionally set the production policies on the "Production" folder
D. * 1 Designate the Organization for production projects only* 2 Ensure that developers do not have the Project Creator IAM role on the Organization * 3 Create development projects outside of the Organization using the developer Google Workspace accounts* 4 Set the policies for all projects on the Organization* 5 Additionally set the production policies on the individual production projects

Answer: C

Explanation:

This option can help create an organization structure that allows developers to create projects, but prevents them from modifying production projects. Folders are containers for projects and other folders within Google Cloud organizations. Folders allow resources to be structured hierarchically and inherit policies from their parent resources. By creating folders under the organization resource named “Development” and “Production”, you can organize your projects by environment and apply different policies to them. By granting all developers the Project Creator IAM role on the “Development” folder, you can allow them to create projects under that folder, but not under the “Production” folder. By moving the developer projects into the “Development” folder, you can ensure that they are subject to the policies set on that folder. By setting the policies for all projects on the organization, you can manage policies centrally and efficiently. By additionally setting the production policies on the “Production” folder, you can enforce more restrictive policies for production projects and prevent developers from modifying them. The other options are not optimal for this scenario, because they either create a second Google Workspace account and organization, which increases complexity and cost (A), or do not use folders to organize projects by environment, which makes it harder to manage policies and permissions (B, D). References:
? https://cloud.google.com/resource-manager/docs/creating-managing-folders
? https://cloud.google.com/architecture/framework/system-design

NEW QUESTION 50
- (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 51
- (Topic 5)
Your organization wants to control IAM policies for different departments independently, but centrally.
Which approach should you take?

A. Multiple Organizations with multiple Folders
B. Multiple Organizations, one for each department
C. A single Organization with Folder for each department
D. A single Organization with multiple projects, each with a central owner

Answer: C

Explanation:
Folders are nodes in the Cloud Platform Resource Hierarchy. A folder can contain projects, other folders, or a combination of both. You can use folders to group projects under an organization in a hierarchy. For example, your organization might contain multiple departments, each with its own set of GCP resources. Folders allow you to group these
resources on a per-department basis. Folders are used to group resources that share common IAM policies. While a folder can contain multiple folders or resources, a given folder or resource can have exactly one parent.
References: https://cloud.google.com/resource-manager/docs/creating-managing-folders

NEW QUESTION 52
- (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 53
- (Topic 5)
Your company has just recently activated Cloud Identity to manage users. The Google Cloud Organization has been configured as wed. The security learn needs to secure protects that will be part of the Organization. They want to prohibit IAM users outside the domain from gaining permissions from now on. What should they do?

A. Configure an organization policy to restrict identities by domain
B. Configure an organization policy to block creation of service accounts

C. Configure Cloud Scheduler o trigger a Cloud Function every hour that removes all users that don't belong to the Cloud identity domain from all projects.
D. Create a technical user (e g . crawler@yourdomain com), and give it the protect owner rote at root organization level Write a bash script that• Lists all me IAM rules of all projects within the organization• Deletes all users that do not belong to the company domainCreate a Compute Engine instance m a project within the Organization and configure gcloud to be executed with technical user credentials Configure a cron job that executes the bash script every hour.

Answer: A

Explanation:
https://cloud.google.com/resource-manager/docs/organization- policy/restricting-domains

NEW QUESTION 58
- (Topic 5)
You are helping the QA team to roll out a new load-testing tool to test the scalability of your primary cloud services that run on Google Compute Engine with Cloud Bigtable. Which three requirements should they include? Choose 3 answers

A. Ensure that the load tests validate the performance of Cloud Bigtable.
B. Create a separate Google Cloud project to use for the load-testing environment.
C. Schedule the load-testing tool to regularly run against the production environment.
D. Ensure all third-party systems your services use are capable of handling high load.
E. Instrument the production services to record every transaction for replay by the load- testing tool.
F. Instrument the load-testing tool and the target services with detailed logging and metrics collection.

Answer: ABF

NEW QUESTION 62
- (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 65
- (Topic 5)
You are creating a solution to remove backup files older than 90 days from your backup Cloud Storage bucket. You want to optimize ongoing Cloud Storage spend. What should you do?

A. Write a lifecycle management rule in XML and push it to the bucket with gsutil.
B. Write a lifecycle management rule in JSON and push it to the bucket with gsutil.
C. Schedule a cron script using gsutil is -lr gs://backups/** to find and remove items older than 90 days.
D. Schedule a cron script using gsutil ls -1 gs://backups/** to find and remove items older than 90 days and schedule it with cron.

Answer: B

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

NEW QUESTION 68
- (Topic 5)
You are developing a globally scaled frontend for a legacy streaming backend data API.
This API expects
events in strict chronological order with no repeat data for proper processing.
Which products should you deploy to ensure guaranteed-once FIFO (first-in, first-out) delivery of data?

A. Cloud Pub/Sub alone
B. Cloud Pub/Sub to Cloud DataFlow
C. Cloud Pub/Sub to Stackdriver
D. Cloud Pub/Sub to Cloud SQL

Answer: B

Explanation:
Reference https://cloud.google.com/pubsub/docs/ordering

NEW QUESTION 72
- (Topic 5)
You have deployed an application on Anthos clusters (formerly Anthos GKE). According to the SRE practices at your company you need to be alerted if the request latency is above a certain threshold for a specified amount of time. What should you do?

A. Enable the Cloud Trace API on your project and use Cloud Monitoring Alerts to send an alert based on the Cloud Trace metrics

B. Configure Anthos Config Management on your cluster and create a yaml file that defines the SLO and alerting policy you want to deploy in your cluster
C. Use Cloud Profiler to follow up the request latenc
D. Create a custom metric in Cloud Monitoring based on the results of Cloud Profiler, and create an Alerting Policy in case this metric exceeds the threshold
E. Install Anthos Service Mesh on your cluste
F. Use the Google Cloud Console to define a Service Level Objective (SLO)

Answer: D

Explanation:
https://cloud.google.com/service-mesh/docs/overview https://cloud.google.com/service-mesh/docs/observability/slo-overview

NEW QUESTION 75
- (Topic 5)
You are building a continuous deployment pipeline for a project stored in a Git source repository and want to ensure that code changes can be verified deploying to production. What should you do?

A. Use Spinnaker to deploy builds to production using the red/black deployment strategy so that changes can easily be rolled back.
B. Use Spinnaker to deploy builds to production and run tests on production deployments.
C. Use Jenkins to build the staging branches and the master branc
D. Build and deploy changes to production for 10% of users before doing a complete rollout.
E. Use Jenkins to monitor tags in the repositor
F. Deploy staging tags to a staging environment for testing.After testing, tag the repository for production and deploy that to the production environment.

Answer: D

Explanation:
Reference: https://github.com/GoogleCloudPlatform/continuous-deployment-on- kubernetes/blob/master/README.md

NEW QUESTION 77
- (Topic 5)
You want to optimize the performance of an accurate, real-time, weather-charting application. The data comes from 50,000 sensors sending 10 readings a second, in the format of a timestamp and sensor reading. Where should you store the data?

A. Google BigQuery
B. Google Cloud SQL
C. Google Cloud Bigtable
D. Google Cloud Storage

Answer: C

Explanation:
It is time-series data, So Big Table. https://cloud.google.com/bigtable/docs/schema-design-time-series
Google Cloud Bigtable is a scalable, fully-managed NoSQL wide-column database that is suitable for both real-time access and analytics workloads.
Good for:
? Low-latency read/write access
? High-throughput analytics
? Native time series support
? Common workloads:
? IoT, finance, adtech
? Personalization, recommendations
? Monitoring
? Geospatial datasets
? Graphs
References: https://cloud.google.com/storage-options/

NEW QUESTION 82
- (Topic 5)
Your solution is producing performance bugs in production that you did not see in staging and test environments. You want to adjust your test and deployment procedures to avoid this problem in the future. What should you do?

A. Deploy fewer changes to production.
B. Deploy smaller changes to production.
C. Increase the load on your test and staging environments.
D. Deploy changes to a small subset of users before rolling out to production.

Answer: C

NEW QUESTION 83
- (Topic 5)
You are using Cloud CDN to deliver static HTTP(S) website content hosted on a Compute Engine instance group. You want to improve the cache hit ratio. What should you do?

A. Customize the cache keys to omit the protocol from the key.
B. Shorten the expiration time of the cached objects.
C. Make sure the HTTP(S) header “Cache-Region” points to the closest region of your users.
D. Replicate the static content in a Cloud Storage bucke
E. Point CloudCDN toward a load balancer on thatbucket.

Answer: A

Explanation:
Reference https://cloud.google.com/cdn/docs/bestpractices#using_custom_cache_keys_to_improve_cache_hit_ratio

NEW QUESTION 84
- (Topic 5)
Your company has decided to make a major revision of their API in order to create better experiences for their developers. They need to keep the old version of the API available and deployable, while allowing new customers and testers to try out the new API. They want to keep the same SSL and DNS records in place to serve both APIs. What should they do?

A. Configure a new load balancer for the new version of the API.
B. Reconfigure old clients to use a new endpoint for the new API.
C. Have the old API forward traffic to the new API based on the path.
D. Use separate backend pools for each API path behind the load balancer.

Answer: D

Explanation:
https://cloud.google.com/endpoints/docs/openapi/lifecycle-management

NEW QUESTION 89
- (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 94
- (Topic 5)
Your company is using Google Cloud. You have two folders under the Organization: Finance and Shopping. The members of the development team are in a Google Group. The development team group has been assigned the Project Owner role on the Organization. You want to prevent the development team from creating resources in projects in the Finance folder. What should you do?

A. Assign the development team group the Project Viewer role on the Finance folder, and assign the development team group the Project Owner role on the Shopping folder.
B. Assign the development team group only the Project Viewer role on the Finance folder.
C. Assign the development team group the Project Owner role on the Shopping folder, and remove the development team group Project Owner role from the Organization.
D. Assign the development team group only the Project Owner role on the Shopping folder.

Answer: C

Explanation:
https://cloud.google.com/resource-manager/docs/cloud-platform-resource- hierarchy
"Roles are always inherited, and there is no way to explicitly remove a permission for a lower-level resource that is granted at a higher level in the resource hierarchy. Given the above example, even if you were to remove the Project Editor role from Bob on the "Test GCP Project", he would still inherit that role from the "Dept Y" folder, so he would still have the permissions for that role on "Test GCP Project"."
Reference: https://cloud.google.com/resource-manager/docs/creating-managing-folders

NEW QUESTION 98
- (Topic 5)
A lead software engineer tells you that his new application design uses websockets and HTTP sessions that are not distributed across the web servers. You want to help him ensure his application will run property on Google Cloud Platform. What should you do?

A. Help the engineer to convert his websocket code to use HTTP streaming.
B. Review the encryption requirements for websocket connections with the security team.
C. Meet with the cloud operations team and the engineer to discuss load balancer options.
D. Help the engineer redesign the application to use a distributed user session service that does not rely on websockets and HTTP sessions.

Answer: C

Explanation:
Google Cloud Platform (GCP) HTTP(S) load balancing provides global load balancing for HTTP(S) requests destined for your instances.
The HTTP(S) load balancer has native support for the WebSocket protocol.

NEW QUESTION 99
- (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 102
- (Topic 5)
To reduce costs, the Director of Engineering has required all developers to move their development infrastructure resources from on-premises virtual machines (VMs) to Google Cloud Platform. These resources go through multiple start/stop events during the day and require state to persist. You have been asked to design the process of running a development environment in Google Cloud while providing cost visibility to the finance department. Which two steps should you take?
Choose 2 answers

A. Use the --no-auto-delete flag on all persistent disks and stop the VM.
B. Use the -auto-delete flag on all persistent disks and terminate the VM.
C. Apply VM CPU utilization label and include it in the BigQuery billing export.
D. Use Google BigQuery billing export and labels to associate cost to groups.
E. Store all state into local SSD, snapshot the persistent disks, and terminate the VM.
F. Store all state in Google Cloud Storage, snapshot the persistent disks, and terminate the VM.

Answer: AD

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

NEW QUESTION 107
- (Topic 5)
All compute Engine instances in your VPC should be able to connect to an Active Directory server on specific ports. Any other traffic emerging from your instances is not allowed. You want to enforce this using VPC firewall rules.
How should you configure the firewall rules?

A. Create an egress rule with priority 1000 to deny all traffic for all instance
B. Create another egress rule with priority 100 to allow the Active Directory traffic for all instances.
C. Create an egress rule with priority 100 to deny all traffic for all instance
D. Create another egress rule with priority 1000 to allow the Active Directory traffic for all instances.
E. Create an egress rule with priority 1000 to allow the Active Directory traffi
F. Rely on the implied denyegress rule with priority 100 to block all traffic for all instances.
G. Create an egress rule with priority 100 to allow the Active Directory traffi
H. Rely on the implied deny egress rule with priority 1000 to block all traffic for all instances.

Answer: B

Explanation:
https://cloud.google.com/vpc/docs/firewalls

NEW QUESTION 110
- (Topic 5)
You want to enable your running Google Container Engine cluster to scale as demand for your application changes.
What should you do?

A. Add additional nodes to your Container Engine cluster using the following command: gcloud container clusters resize CLUSTER_NAME --size 10
B. Add a tag to the instances in the cluster with the following command:gcloud compute instances add-tags INSTANCE --tags enable --autoscaling max-nodes-10
C. Update the existing Container Engine cluster with the following command:gcloud alpha container clusters update mycluster --enable-autoscaling --min-nodes=1 -- max-nodes=10
D. Create a new Container Engine cluster with the following command:gcloud alpha container clusters create mycluster --enable-autocaling --min-nodes=1 --max-nodes=10and redeploy your application.

Answer: B

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/cluster- autoscaler
Cluster autoscaling
--enable-autoscaling
Enables autoscaling for a node pool.
Enables autoscaling in the node pool specified by --node-pool or the default node pool if -- node-pool is not provided.
Where:
--max-nodes=MAX_NODES
Maximum number of nodes in the node pool.
Maximum number of nodes to which the node pool specified by --node-pool (or default node pool if unspecified) can scale.

NEW QUESTION 113
- (Topic 5)

You are creating an App Engine application that uses Cloud Datastore as its persistence layer. You need to retrieve several root entities for which you have the identifiers. You want to minimize the overhead in operations performed by Cloud Datastore. What should you do?

A. Create the Key object for each Entity and run a batch get operation
B. Create the Key object for each Entity and run multiple get operations, one operation for each entity
C. Use the identifiers to create a query filter and run a batch query operation
D. Use the identifiers to create a query filter and run multiple query operations, one operation for each entity

Answer: C

Explanation:
https://cloud.google.com/datastore/docs/concepts/entities#datastore-datastore-batch-upsert-nodejs

NEW QUESTION 118
- (Topic 5)
The development team has provided you with a Kubernetes Deployment file. You have no infrastructure yet and need to deploy the application. What should you do?

A. Use gcloud to create a Kubernetes cluste
B. Use Deployment Manager to create the deployment.
C. Use gcloud to create a Kubernetes cluste
D. Use kubect1 to create the deployment.
E. Use kubect1 to create a Kubernetes cluste
F. Use Deployment Manager to create the deployment.
G. Use kubect1 to create a Kubernetes cluste
H. Use kubect1 to create the deployment.

Answer: B

Explanation:
https://cloud.google.com/kubernetes-engine/docs/how-to/creating-a-cluster

NEW QUESTION 121
- (Topic 5)
You have an application that runs in Google Kubernetes Engine (GKE). Over the last 2 weeks, customers have reported that a specific part of the application returns errors very frequently. You currently have no logging or monitoring solution enabled on your GKE cluster. You want to diagnose the problem, but you have not been able to replicate the issue. You want to cause minimal disruption to the application. What should you do?

A. * 1. Update your GKE cluster to use Cloud Operations for GKE.* 2. Use the GKE Monitoring dashboard to investigate logs from affected Pods.
B. * 1. Create a new GKE cluster with Cloud Operations for GKE enabled.* 2. Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluster.* 3. Use the GKE Monitoring dashboard to investigate logs from affected Pods.
C. * 1. Update your GKE cluster to use Cloud Operations for GKE, and deploy Prometheus.* 2. Set an alert to trigger whenever the application returns an error.
D. * 1. Create a new GKE cluster with Cloud Operations for GKE enabled, and deploy Prometheus.* 2. Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluster.* 3. Set an alert to trigger whenever the application returns an error.

Answer: A

Explanation:
Reference: https://cloud.google.com/blog/products/management-tools/using-logging-your- apps-running- kubernetes-engine

NEW QUESTION 126
- (Topic 5)
Your company is migrating its on-premises data center into the cloud. As part of the migration, you want to integrate Kubernetes Engine for workload orchestration. Parts of your architecture must also be PCI DSScompliant.
Which of the following is most accurate?

A. App Engine is the only compute platform on GCP that is certified for PCI DSS hosting.
B. Kubernetes Engine cannot be used under PCI DSS because it is considered shared hosting.
C. Kubernetes Engine and GCP provide the tools you need to build a PCI DSS-compliant environment.
D. All Google Cloud services are usable because Google Cloud Platform is certified PCI- compliant.

Answer: D

Explanation:
https://cloud.google.com/security/compliance/pci-dss

NEW QUESTION 130
- (Topic 5)
You deploy your custom java application to google app engine. It fails to deploy and gives you the following stack trace:

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

Answer: B

NEW QUESTION 135
- (Topic 5)
You are deploying an application to Google Cloud. The application is part of a system. The application in Google Cloud must communicate over a private network with applications in a non-Google Cloud environment. The expected average throughput is 200 kbps. The business requires:
• 99.99% system availability
• cost optimization
You need to design the connectivity between the locations to meet the business requirements. What should you provision?

A. A Classic Cloud VPN gateway connected with one tunnel to an on-premises VPN gateway.
B. A Classic Cloud VPN gateway connected with two tunnels to an on-premises VPN gateway.
C. An HA Cloud VPN gateway connected with two tunnels to an on-premises VPN gateway.
D. Two HA Cloud VPN gateways connected to two on-premises VPN gateway
E. Configure each HA CloudVPN gateway to have two tunnels, each connected to different on-premises VPN gateways.

Answer: C

Explanation:
https://cloud.google.com/network-connectivity/docs/vpn/concepts/topologies#configurations_that_support_9999_availability

NEW QUESTION 139
- (Topic 5)
Your team will start developing a new application using microservices architecture on Kubernetes Engine. As part of the development lifecycle, any code change that has been pushed to the remote develop branch on your GitHub repository should be built and tested automatically. When the build and test are successful, the relevant microservice will be deployed automatically in the development environment. You want to ensure that all code deployed in the development environment follows this process. What should you do?

A. Have each developer install a pre-commit hook on their workstation that tests the code and builds the container when committing on the development branc
B. After a successful commit, have the developer deploy the newly built container image on the development cluster.
C. Install a post-commit hook on the remote git repository that tests the code and builds the container when code is pushed to the development branc
D. After a successful commit, have the developer deploy the newly built container image on the development cluster.
E. Create a Cloud Build trigger based on the development branch that tests the code, builds the container, and stores it in Container Registr
F. Create a deployment pipeline that watches for new images and deploys the new image on the development cluste
G. Ensure only the deployment tool has access to deploy new versions.
H. Create a Cloud Build trigger based on the development branch to build a new container image and store it in Container Registr

I. Rely on Vulnerability Scanning to ensure the code tests succee
J. As the final step of the Cloud Build process, deploy the new container image on the development cluste
K. Ensure only Cloud Build has access to deploy new versions.

Answer: C

Explanation:
https://cloud.google.com/container-registry/docs/overview
Create a Cloud Build trigger based on the development branch that tests the code, builds the container, and stores it in Container Registry. Create a deployment pipeline that watches for new images and deploys the new image on the development cluster. Ensure only the deployment tool has access to deploy new versions.

NEW QUESTION 144
- (Topic 5)
Your company uses Google Kubernetes Engine (GKE) as a platform for all workloads. Your company has a single large GKE cluster that contains batch, stateful, and stateless workloads. The GKE cluster is configured with a single node pool with 200 nodes. Your company needs to reduce the cost of this cluster but does not want to compromise availability. What should you do?

A. Create a second GKE cluster for the batch workloads onl
B. Allocate the 200 original nodes across both clusters.
C. Configure a HorizontalPodAutoscaler for all stateless workloads and for all compatible stateful workload
D. Configure the cluster to use node auto scaling.
E. Configure CPU and memory limits on the namespaces in the cluste
F. Configure all Pods to have a CPU and memory limits.
G. Change the node pool to use spot VMs.

Answer: C

Explanation:
One way to reduce the cost of a Google Kubernetes Engine (GKE) cluster without compromising availability is to use horizontal pod autoscalers (HPA) and node auto scaling. HPA allows you to automatically scale the number of Pods in a deployment based on the resource usage of the Pods. By configuring HPA for stateless workloads and for compatible stateful workloads, you can ensure that the number of Pods is automatically adjusted based on the actual resource usage, which can help to reduce costs. Node auto scaling allows you to automatically add or remove nodes from the node pool based on the resource usage of the cluster. By configuring node auto scaling, you can ensure that the cluster has the minimum number of nodes needed to meet the resource requirements of the workloads, which can also help to reduce costs.

NEW QUESTION 147
- (Topic 5)
You have deployed an application to Kubernetes Engine, and are using the Cloud SQL proxy container to make the Cloud SQL database available to the services running on Kubernetes. You are notified that the application is reporting database connection issues. Your company policies require a post- mortem. What should you do?

A. Use gcloud sql instances restart.
B. Validate that the Service Account used by the Cloud SQL proxy container still has the Cloud Build Editor role.
C. In the GCP Console, navigate to Stackdriver Loggin
D. Consult logs for Kubernetes Engine and Cloud SQL.
E. In the GCP Console, navigate to Cloud SQ
F. Restore the latest backu
G. Use kubect1 to restart all pods.

Answer: C

NEW QUESTION 149
- (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 152
- (Topic 5)
Google Cloud Platform resources are managed hierarchically using organization, folders, and projects. When Cloud Identity and Access Management (IAM) policies exist at these different levels, what is the effective policy at a particular node of the hierarchy?

A. The effective policy is determined only by the policy set at the node
B. The effective policy is the policy set at the node and restricted by the policies of its ancestors
C. The effective policy is the union of the policy set at the node and policies inherited from its ancestors
D. The effective policy is the intersection of the policy set at the node and policies inherited from its ancestors

Answer: B

Explanation:
Reference: https://cloud.google.com/resource-manager/docs/cloud-platform-resource- hierarchy

NEW QUESTION 156
- (Topic 5)
Your company is building a new architecture to support its data-centric business focus. You are responsible for setting up the network. Your company's mobile and web-facing applications will be deployed on-premises, and all data analysis will be conducted in GCP. The plan is to process and load 7 years of archived .csv files totaling 900 TB of data and then continue loading 10 TB of data daily. You currently have an existing 100-MB internet connection.
What actions will meet your company's needs?

A. Compress and upload both achieved files and files uploaded daily using the qsutil –m option.
B. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transferarchived data to Cloud Storag
C. Establish a connection with Google using a Dedicated Interconnect orDirect Peering connection and use it to upload files daily.
D. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transferarchived data to Cloud Storag
E. Establish one Cloud VPN Tunnel to VPC networks over the public internet, and compares and upload files daily using the gsutil –m option.
F. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storag
G. Establish a Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily.

Answer: B

Explanation:
https://cloud.google.com/interconnect/docs/how-to/direct-peering

NEW QUESTION 157
- (Topic 5)
You are working at a financial institution that stores mortgage loan approval documents on Cloud Storage. Any change to these approval documents must be uploaded as a separate approval file, so you want to ensure that these documents cannot be deleted or overwritten for the next 5 years. What should you do?

A. Create a retention policy on the bucket for the duration of 5 year
B. Create a lock on the retention policy.
C. Create the bucket with uniform bucket-level access, and grant a service account the role of Object Write
D. Use the service account to upload new files.
E. Use a customer-managed key for the encryption of the bucke
F. Rotate the key after 5 years.
G. Create the bucket with fine-grained access control, and grant a service account the role of Object Write
H. Use the service account to upload new files.

Answer: A

Explanation:
Reference: https://cloud.google.com/storage/docs/using-bucket-lock

NEW QUESTION 161
- (Topic 5)
You have an outage in your Compute Engine managed instance group: all instance keep restarting after 5 seconds. You have a health check configured, but autoscaling is disabled. Your colleague, who is a Linux expert, offered to look into the issue. You need to make sure that he can access the VMs. What should you do?

A. Grant your colleague the IAM role of project Viewer
B. Perform a rolling restart on the instance group
C. Disable the health check for the instance grou
D. Add his SSH key to the project-wide SSH keys
E. Disable autoscaling for the instance grou
F. Add his SSH key to the project-wide SSH Keys

Answer: C

Explanation:
https://cloud.google.com/compute/docs/instance-groups/autohealing-instances-in-migs Health checks used for autohealing should be conservative so they don't preemptively delete and recreate your instances. When an autohealer health check is too aggressive, the autohealer might mistake busy instances for failed instances and unnecessarily restart them, reducing availability

NEW QUESTION 164
- (Topic 5)
Your company wants to track whether someone is present in a meeting room reserved for a scheduled meeting. There are 1000 meeting rooms across 5 offices on 3 continents. Each room is equipped with a motion sensor that reports its status every second. The data from the motion detector includes only a sensor ID and several different discrete items of information. Analysts will use this data, together with information about account owners and office locations. Which database type should you use?

A. Flat file
B. NoSQL
C. Relational
D. Blobstore

Answer: B

Explanation:
Relational databases were not designed to cope with the scale and agility challenges that face modern applications, nor were they built to take advantage of the commodity storage and processing power available today.
NoSQL fits well for:
_ Developers are working with applications that create massive volumes of new, rapidly changing data types — structured, semi-structured, unstructured and polymorphic data.

NEW QUESTION 166
- (Topic 5)
An application development team has come to you for advice.They are planning to write and deploy an HTTP(S) API using Go 1.12. The API will have a very unpredictable workload and must remain reliable during peaks in traffic. They want to minimize operational overhead for this application. What approach should you recommend?

A. Use a Managed Instance Group when deploying to Compute Engine
B. Develop an application with containers, and deploy to Google Kubernetes Engine (GKE)
C. Develop the application for App Engine standard environment
D. Develop the application for App Engine Flexible environment using a custom runtime

Answer: C

Explanation:
https://cloud.google.com/appengine/docs/the-appengine-environments

NEW QUESTION 169
- (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 172
- (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 175
- (Topic 5)
Your company has a stateless web API that performs scientific calculations. The web API runs on a single Google Kubernetes Engine (GKE) cluster. The cluster is currently deployed in us-central1. Your company has expanded to offer your API to customers in Asia. You want to reduce the latency for the users in Asia. What should you do?

A. Use a global HTTP(s) load balancer with Cloud CDN enabled
B. Create a second GKE cluster in asia-southeast1, and expose both API's using a Service oftype Load Balance
C. Add the public Ips to the Cloud DNS zone
D. Increase the memory and CPU allocated to the application in the cluster
E. Create a second GKE cluster in asia-southeast1, and use kubemci to create a global HTTP(s) load balancer

Answer: D

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/multi-cluster- ingress#how_works
https://github.com/GoogleCloudPlatform/k8s-multicluster-ingress https://cloud.google.com/blog/products/gcp/how-to-deploy-geographically-distributed-services-on-kubernetes-engine-with-kubemci

NEW QUESTION 178
- (Topic 5)
You need to deploy an application to Google Cloud. The application receives traffic via TCP and reads and writes data to the filesystem. The application does not support horizontal scaling. The application process requires full control over the data on the file system because concurrent access causes corruption. The business is willing to accept a downtime when an incident occurs, but the application must be available 24/7 to support their business operations. You need to design the architecture of this application on Google Cloud.
What should you do?

A. Use a managed instance group with instances in multiple zones, use Cloud Filestore, and use an HTTP load balancer in front of the instances.
B. Use a managed instance group with instances in multiple zones, use Cloud Filestore, and use a network load balancer in front of the instances.
C. Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use an HTTP load balancer in front of the instances.
D. Use an unmanaged instance group with an active and standby instance in different zones, use a regional persistent disk, and use a network load balancer in

front of the instances.

Answer: D

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

NEW QUESTION 182
- (Topic 5)
You are designing an application for use only during business hours. For the minimum viable product release, you'd like to use a managed product that automatically "scales to zero" so you don't incur costs when there is no activity.
Which primary compute resource should you choose?

A. Cloud Functions
B. Compute Engine
C. Kubernetes Engine
D. AppEngine flexible environment

Answer: A

Explanation:
https://cloud.google.com/serverless-options

NEW QUESTION 183
- (Topic 5)
Your web application must comply with the requirements of the European Union's General Data Protection Regulation (GDPR). You are responsible for the technical architecture of your web application. What should you do?

A. Ensure that your web application only uses native features and services of Google Cloud Platform,because Google already has various certifications and provides "pass-on" compliance when you use native features.
B. Enable the relevant GDPR compliance setting within the GCPConsole for each of the services in use within your application.
C. Ensure that Cloud Security Scanner is part of your test planning strategy in order to pick up any compliance gaps.
D. Define a design for the security of data in your web application that meets GDPRrequirements.

Answer: D

Explanation:
https://cloud.google.com/security/gdpr/?tab=tab4
Reference: https://www.mobiloud.com/blog/gdpr-compliant-mobile-app/

NEW QUESTION 185
- (Topic 5)
For this question, refer to the TerramEarth case study. You are building a microservice- based application for TerramEarth. The application is based on Docker containers. You want to follow Google-recommended practices to build the application continuously and store the build artifacts. What should you do?
A.
* 1. Configure a trigger in Cloud Build for new source changes.
* 2. Invoke Cloud Build to build one container image, and tag the image with the label 'latest.'
* 3. Push the image to the Artifact Registry.
B.
* 1. Configure a trigger in Cloud Build for new source changes.
* 2. Invoke Cloud Build to build container images for each microservice, and tag them using the code commit hash.
* 3. Push the images to the Artifact Registry.
C.
* 1 Create a Scheduler job to check the repo every minute.
* 2. For any new change, invoke Cloud Build to build container images for the microservices.
* 3. Tag the images using the current timestamp, and push them to the Artifact Registry.
D.
* 1. Configure a trigger in Cloud Build for new source changes.
* 2. The trigger invokes build jobs and build container images for the microservices.
* 3. Tag the images with a version number, and push them to Cloud Storage.

A.

Answer: C

NEW QUESTION 189
- (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 192
- (Topic 5)
You want to automate the creation of a managed instance group and a startup script to install the OS package dependencies. You want to minimize the startup time for VMs in the instance group.
What should you do?

A. Use Terraform to create the managed instance group and a startup script to install the OS packagedependencies.
B. Create a custom VM image with all OS package dependencie
C. Use Deployment Manager to create the managed instance group with the VM image.
D. Use Puppet to create the managed instance group and install the OS package dependencies.
E. Use Deployment Manager to create the managed instance group and Ansible to install the OS package dependencies.

Answer: B

Explanation:
"Custom images are more deterministic and start more quickly than instances with startup scripts. However, startup scripts are more flexible and let you update the apps and settings in your instances more easily." https://cloud.google.com/compute/docs/instance- templates/create-instance-templates#using_custom_or_public_images_in_your_instance_templates

NEW QUESTION 193
- (Topic 5)
You have been engaged by your client to lead the migration of their application infrastructure to GCP. One of their current problems is that the on-premises high performance SAN is requiring frequent and expensive upgrades to keep up with the variety of workloads that are identified as follows: 20TB of log archives retained for legal reasons; 500 GB of VM boot/data volumes and templates; 500 GB of image thumbnails; 200 GB of customer session state data that allows customers to restart sessions even if off-line for several days.
Which of the following best reflects your recommendations for a cost-effective storage allocation?

A. Local SSD for customer session state dat
B. Lifecycle-managed Cloud Storage for logarchives, thumbnails, and VM boot/data volumes.
C. Memcache backed by Cloud Datastore for the customer session state dat
D. Lifecycle- managed CloudStorage for log archives, thumbnails, and VM boot/data volumes.
E. Memcache backed by Cloud SQL for customer session state dat
F. Assorted local SSD- backed instances for VM boot/data volume
G. Cloud Storage for log archives and thumbnails.
H. Memcache backed by Persistent Disk SSD storage for customer session state dat
I. Assorted local SSDbacked instances for VM boot/data volume
J. Cloud Storage for log archives and thumbnails.

Answer: D

Explanation:
https://cloud.google.com/compute/docs/disks

NEW QUESTION 195
- (Topic 5)
You created a pipeline that can deploy your source code changes to your infrastructure in instance groups for self healing.
One of the changes negatively affects your key performance indicator. You are not sure how to fix it and investigation could take up to a week. What should you do

A. Log in to a server, and iterate a fix locally
B. Change the instance group template to the previous one, and delete all instances.
C. Revert the source code change and rerun the deployment pipeline
D. Log into the servers with the bad code change, and swap in the previous code

Answer: C

NEW QUESTION 200
- (Topic 5)
An application development team believes their current logging tool will not meet their needs for their new cloud-based product. They want a bettor tool to capture errors and help them analyze their historical log data. You want to help them find a solution that meets their needs, what should you do?

A. Direct them to download and install the Google StackDriver logging agent.
B. Send them a list of online resources about logging best practices.
C. Help them define their requirements and assess viable logging tools.
D. Help them upgrade their current tool to take advantage of any new features.

Answer: C

Explanation:
Help them define their requirements and assess viable logging tools. They know the requirements and the existing tools' problems. While it's true StackDriver Logging and Error Reporting possibly meet all their requirements, there might be other tools also meet their need. They need you to provide expertise to make assessment for new tools, specifically, logging tools that can "capture errors and help them analyze their historical log data".
References: https://cloud.google.com/logging/docs/agent/installation

NEW QUESTION 202
- (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 206
- (Topic 5)
You need to develop procedures to verify resilience of disaster recovery for remote recovery using GCP. Your production environment is hosted on-premises. You need to establish a secure, redundant connection between your on premises network and the GCP network.
What should you do?

A. Verify that Dedicated Interconnect can replicate files to GC
B. Verify that direct peering can establish asecure connection between your networks if Dedicated Interconnect fails.
C. Verify that Dedicated Interconnect can replicate files to GC
D. Verify that Cloud VPN can establish a secure connection between your networks if Dedicated Interconnect fails.
E. Verify that the Transfer Appliance can replicate files to GC
F. Verify that direct peering can establish asecure connection between your networks if the Transfer Appliance fails.
G. Verify that the Transfer Appliance can replicate files to GC
H. Verify that Cloud VPN can establish a secure connection between your networks if the Transfer Appliance fails.

Answer: B

Explanation:
https://cloud.google.com/interconnect/docs/how-to/direct-peering

NEW QUESTION 211
- (Topic 5)
A recent audit that a new network was created in Your GCP project. In this network, a GCE instance has an SSH port open the world. You want to discover this network's origin. What should you do?

A. Search for Create VM entry in the Stackdriver alerting console.
B. Navigate to the Activity page in the Home sectio
C. Set category to Data Access and search for Create VM entry.
D. In the logging section of the console, specify GCE Network as the logging sectio
E. Search for the Create Insert entry.
F. Connect to the GCE instance using project SSH Key
G. Identify previous logins in system logs, and match these with the project owners list.

Answer: C

NEW QUESTION 214
- (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 216
- (Topic 5)
Your company has an enterprise application running on Compute Engine that requires high availability and high performance. The application has been deployed on two instances in two zones in the same region m active passive mode. The application writes data to a persistent disk in the case of a single zone outage that data should be immediately made available to the other instance in the other zone. You want to maximize performance while minimizing downtime and data loss.
What should you do?
A.
* 1. Attach a persistent SSD disk to the first instance
* 2. Create a snapshot every hour
* 3. In case of a zone outage, recreate a persistent SSD disk in the second instance where data is coming from the created snapshot
B.
* 1 Create a Cloud Storage bucket
* 2. Mount the bucket into the first instance with gcs-fuse
* 3. In case of a zone outage, mount the Cloud Storage bucket to the second instance with gcs-fuse
C.

* 1 Attach a local SSD lo the first instance disk
* 2. Execute an rsync command every hour where the target is a persistent SSD disk attached to the second instance
* 3. In case of a zone outage, use the second instance
D.
* 1. Attach a regional SSD persistent Ask to the first instance
* 2. In case of a zone outage, force-attach the disk to the other instance

A.

Answer: D

NEW QUESTION 221
- (Topic 5)
Your company has multiple on-premises systems that serve as sources for reporting. The data has not been maintained well and has become degraded over time. You want to use Google-recommended practices to detect anomalies in your company data. What should you do?

A. Upload your files into Cloud Storag
B. Use Cloud Datalab to explore and clean your data.
C. Upload your files into Cloud Storag
D. Use Cloud Dataprep to explore and clean your data.
E. Connect Cloud Datalab to your on-premises system
F. Use Cloud Datalab to explore and clean your data.
G. Connect Cloud Dataprep to your on-premises system
H. Use Cloud Dataprep to exploreand clean your data.

Answer: B

Explanation:
https://cloud.google.com/dataprep/

NEW QUESTION 223
- (Topic 5)
Your team is developing a web application that will be deployed on Google Kubernetes Engine (GKE). Your CTO expects a successful launch and you need to ensure your application can handle the expected load of tens of thousands of users. You want to test the current deployment to ensure the latency of your application stays below a certain threshold. What should you do?

A. Use a load testing tool to simulate the expected number of concurrent users and total requests to your application, and inspect the results.
B. Enable autoscaling on the GKE cluster and enable horizontal pod autoscaling on your application deployment
C. Send curl requests to your application, and validate if the auto scaling works.
D. Replicate the application over multiple GKE clusters in every Google Cloud region.Configure a global HTTP(S) load balancer to expose the different clusters over a single global IP address.
E. Use Cloud Debugger in the development environment to understand the latency between the different microservices.

Answer: B

NEW QUESTION 227
- (Topic 5)
You are managing an application deployed on Cloud Run for Anthos, and you need to define a strategy for deploying new versions of the application. You want to evaluate the new code with a subset of production traffic to decide whether to proceed with the rollout. What should you do?

A. Deploy a new revision to Cloud Run with the new versio
B. Configure traffic percentage between revisions.
C. Deploy a new service to Cloud Run with the new versio
D. Add a Cloud Load Balancing instance in front of both services.
E. In the Google Cloud Console page for Cloud Run, set up continuous deployment using Cloud Build for the development branc
F. As part of the Cloud Build trigger, configure the substitution variable TRAFFIC_PERCENTAGE with the percentage of traffic you want directed to a new version.
G. In the Google Cloud Console, configure Traffic Director with a new Service that points to the new version of the application on Cloud Ru
H. Configure Traffic Director to send a small percentage of traffic to the new version of the application.

Answer: A

Explanation:
https://cloud.google.com/run/docs/rollouts-rollbacks-traffic-migration

NEW QUESTION 228
- (Topic 5)
Your web application uses Google Kubernetes Engine to manage several workloads. One workload requires a consistent set of hostnames even after pod scaling and relaunches.
Which feature of Kubernetes should you use to accomplish this?

A. StatefulSets
B. Role-based access control
C. Container environment variables
D. Persistent Volumes

Answer: A

Explanation:
https://kubernetes.io/docs/tutorials/stateful-application/basic-stateful-set/

NEW QUESTION 229
- (Topic 5)
Your company has an application running on App Engine that allows users to upload music files and share them with other people. You want to allow users to upload files directly into Cloud Storage from their browser session. The payload should not be passed through the backend. What should you do?
A.
* 1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App Engine application is an allowed origin.
* 2. Use the Cloud Storage Signed URL feature to generate a POST URL.
B.
* 1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App Engine application is an allowed origin.
* 2. Assign the Cloud Storage WRITER role to users who upload files.
C.
* 1. Use the Cloud Storage Signed URL feature to generate a POST URL.
* 2. Use App Engine default credentials to sign requests against Cloud Storage.
D.
* 1. Assign the Cloud Storage WRITER role to users who upload files.
* 2. Use App Engine default credentials to sign requests against Cloud Storage.

A.

Answer: B

NEW QUESTION 234
- (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 235
- (Topic 5)
Your company uses the Firewall Insights feature in the Google Network Intelligence Center. You have several firewall rules applied to Compute Engine instances. You need to evaluate the efficiency of the applied firewall ruleset. When you bring up the Firewall Insights page in the Google Cloud Console, you notice that there are no log rows to display. What should you do to troubleshoot the issue?

A. Enable Virtual Private Cloud (VPC) flow logging.
B. Enable Firewall Rules Logging for the firewall rules you want to monitor.
C. Verify that your user account is assigned the compute.networkAdmin Identity and Access Management (IAM) role.
D. Install the Google Cloud SDK, and verify that there are no Firewall logs in the command line output.

Answer: B

Explanation:

Reference: https://cloud.google.com/network-intelligence-center/docs/firewall-insights/how-to/using-firewall- insights

NEW QUESTION 240
- (Topic 5)
Your company is forecasting a sharp increase in the number and size of Apache Spark and Hadoop jobs being run on your local datacenter You want to utilize the cloud to help you scale this upcoming demand with the least amount of operations work and code change. Which product should you use?

A. Google Cloud Dataflow
B. Google Cloud Dataproc
C. Google Compute Engine
D. Google Container Engine

Answer: B

Explanation:
Google Cloud Dataproc is a fast, easy-to-use, low-cost and fully managed service that lets you run the Apache Spark and Apache Hadoop ecosystem on Google Cloud Platform. Cloud Dataproc provisions big or small clusters rapidly, supports many popular job types, and is integrated with other Google Cloud Platform services, such as Google Cloud Storage and Stackdriver Logging, thus helping you reduce TCO.
References: https://cloud.google.com/dataproc/docs/resources/faq

NEW QUESTION 243
- (Topic 5)
The operations manager asks you for a list of recommended practices that she should consider when migrating a J2EE application to the cloud. Which three practices should you recommend? Choose 3 answers

A. Port the application code to run on Google App Engine.
B. Integrate Cloud Dataflow into the application to capture real-time metrics.
C. Instrument the application with a monitoring tool like Stackdriver Debugger.
D. Select an automation framework to reliably provision the cloud infrastructure.
E. Deploy a continuous integration tool with automated testing in a staging environment.
F. Migrate from MySQL to a managed NoSQL database like Google Cloud Datastore or Bigtable.

Answer: AEF

Explanation:
References: https://cloud.google.com/appengine/docs/standard/java/tools/uploadinganapp https://cloud.google.com/appengine/docs/standard/java/building-app/cloud-sql

NEW QUESTION 247
- (Topic 5)
Your company places a high value on being responsive and meeting customer needs quickly. Their primary business objectives are release speed and agility. You want to reduce the chance of security errors being accidentally introduced. Which two actions can you take? Choose 2 answers

A. Ensure every code check-in is peer reviewed by a security SME.
B. Use source code security analyzers as part of the CI/CD pipeline.
C. Ensure you have stubs to unit test all interfaces between components.
D. Enable code signing and a trusted binary repository integrated with your CI/CD pipeline.
E. Run a vulnerability security scanner as part of your continuous-integration /continuous- delivery (CI/CD) pipeline.

Answer: BE

Explanation:
https://docs.microsoft.com/en-us/vsts/articles/security-validation-cicd-pipeline?view=vsts

NEW QUESTION 251
- (Topic 5)
Your company has decided to build a backup replica of their on-premises user authentication PostgreSQL database on Google Cloud Platform. The database is 4 TB, and large updates are frequent. Replication requires private address space communication. Which networking approach should you use?

A. Google Cloud Dedicated Interconnect
B. Google Cloud VPN connected to the data center network
C. A NAT and TLS translation gateway installed on-premises
D. A Google Compute Engine instance with a VPN server installed connected to the data center network

Answer: A

Explanation:
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations
Google Cloud Dedicated Interconnect provides direct physical connections and RFC 1918 communication between your on-premises network and Google's network. Dedicated Interconnect enables you to transfer large amounts of data between networks, which can be more cost effective than purchasing additional bandwidth over the public Internet or using VPN tunnels.
Benefits:
? Traffic between your on-premises network and your VPC network doesn't traverse the public Internet. Traffic traverses a dedicated connection with fewer hops, meaning there are less points of failure where traffic might get dropped or disrupted.
? Your VPC network's internal (RFC 1918) IP addresses are directly accessible from your on-premises network. You don't need to use a NAT device or VPN tunnel to reach internal IP addresses. Currently, you can only reach internal IP addresses over a dedicated connection. To reach Google external IP addresses, you must use a separate connection.
? You can scale your connection to Google based on your needs. Connection capacity is delivered over one or more 10 Gbps Ethernet connections, with a maximum of eight connections (80 Gbps total per interconnect).
? The cost of egress traffic from your VPC network to your on-premises network is reduced. A dedicated connection is generally the least expensive method if you have a high-volume of traffic to and from Google's network.
References: https://cloud.google.com/interconnect/docs/details/dedicated

NEW QUESTION 256
- (Topic 5)
You are migrating third-party applications from optimized on-premises virtual machines to Google Cloud. You are unsure about the optimum CPU and memory options. The application have a consistent usage patterns across multiple weeks. You want to optimize resource usage for the lowest cost. What should you do?

A. Create a Compute engine instance with CPU and Memory options similar to your application's current on-premises virtual machin
B. Install the cloud monitoring agent, and deploy the third party applicatio
C. Run a load with normal traffic levels on third party application and follow the Rightsizing Recommendations in the Cloud Console
D. Create an App Engine flexible environment, and deploy the third party application using a Docker file and a custom runtim
E. Set CPU and memory options similar to your application's current on-premises virtual machine in the app.yaml file.
F. Create an instance template with the smallest available machine type, and use an imageof the third party application taken from the current on-premises virtual machin
G. Create a managed instance group that uses average CPU to autoscale the number of instances in the grou
H. Modify the average CPU utilization threshold to optimize the number of instances running.
I. Create multiple Compute Engine instances with varying CPU and memory option
J. Install the cloud monitoring agent and deploy the third-party application on each of the
K. Run a load test with high traffic levels on the application and use the results to determine the optimal settings.

Answer: A

Explanation:
Create a Compute engine instance with CPU and Memory options similar to your application's current on-premises virtual machine. Install the cloud monitoring agent, and deploy the third party application. Run a load with normal traffic levels on third party application and follow the Rightsizing Recommendations in the Cloud Console https://cloud.google.com/migrate/compute-engine/docs/4.9/concepts/planning-a-migration/cloud-instance-rightsizing?hl=en

NEW QUESTION 260
- (Topic 5)
You have an application that will run on Compute Engine. You need to design an architecture that takes into account a disaster recovery plan that requires your application to fail over to another region in case of a regional outage. What should you do?

A. Deploy the application on two Compute Engine instances in the same project but in a different regio
B. Use the first instance to serve traffic, and use the HTTP load balancing service to fail over to the standby instance in case of a disaster.
C. Deploy the application on a Compute Engine instanc
D. Use the instance to serve traffic, and use the HTTP load balancing service to fail over to an instance on your premises in case of a disaster.
E. Deploy the application on two Compute Engine instance groups, each in the same project but in a different regio
F. Use the first instance group to serve traffic, and use the HTTP load balancing service to fail over to the standby instance group in case of a disaster.
G. Deploy the application on two Compute Engine instance groups, each in separate project and a different regio
H. Use the first instance group to server traffic, and use the HTTP load balancing service to fail over to the standby instance in case of a disaster.

Answer: C

NEW QUESTION 264
- (Topic 5)
You are working at an institution that processes medical data. You are migrating several workloads onto Google Cloud. Company policies require all workloads to run on physically separated hardware, and workloads from different clients must also be separated You created a sole-tenant node group and added a node for each client. You need to deploy the workloads on these dedicated hosts. What should you do?

A. Add the node group name as a network tag when creating Compute Engine instances in order to host each workload on the correct node group.
B. Add the node name as a network tag when creating Compute Engine instances in order to host each workload on the correct node.
C. Use node affinity labels based on the node group name when creating Compute Engine instances in order to host each workload on the correct node group
D. Use node affinity labels based on the node name when creating Compute Engine instances in order to host each workload on the correct node.

Answer: C

Explanation:
https://cloud.google.com/compute/docs/nodes/provisioning-sole-tenant-vms#provision_a_sole-tenant_vm
https://cloud.google.com/compute/docs/nodes/provisioning-sole-tenant-vms#gcloud_2 When you create a VM, you request sole-tenancy by specifying node affinity or anti-affinity, referencing one or more node affinity labels. You specify custom node affinity labels when you create a node template, and Compute Engine automatically includes some default affinity labels on each node. By specifying affinity when you create a VM, you can schedule VMs together on a specific node or nodes in a node group. By specifying anti-affinity when you create a VM, you can ensure that certain VMs are not scheduled together on the same node or nodes in a node group.

NEW QUESTION 266
- (Topic 5)
Your company is designing its application landscape on Compute Engine. Whenever a
zonal outage occurs, the application should be restored in another zone as quickly as possible with the latest application data. You need to design the solution to meet this requirement. What should you do?

A. Create a snapshot schedule for the disk containing the application dat
B. Whenever a zonal outage occurs, use the latest snapshot to restore the disk in the same zone.
C. Configure the Compute Engine instances with an instance template for the application, and use a regional persistent disk for the application dat
D. Whenever a zonal outage occurs, use the instance template to spin up the application in another zone in the same regio
E. Use the regional persistent disk for the application data.
F. Create a snapshot schedule for the disk containing the application dat
G. Whenever a zonal outage occurs, use the latest snapshot to restore the disk in another zone within the same region.
H. Configure the Compute Engine instances with an instance template for the application, and use a regional persistent disk for the application dat
I. Whenever a zonal outage occurs, use the instance template to spin up the application in another regio
J. Use the regional persistent disk for the application data,

Answer: B

Explanation:
Regional persistent disk is a storage option that provides synchronous replication of data between two zones in a region. Regional persistent disks can be a good building block to use when you implement HA services in Compute Engine. https://cloud.google.com/compute/docs/disks/high-availability-regional-persistent-disk

NEW QUESTION 268
- (Topic 5)
Your customer is moving an existing corporate application to Google Cloud Platform from an on-premises data center. The business owners require minimal user disruption. There are strict security team requirements for storing passwords. What authentication strategy should they use?

A. Use G Suite Password Sync to replicate passwords into Google.
B. Federate authentication via SAML 2.0 to the existing Identity Provider.
C. Provision users in Google using the Google Cloud Directory Sync tool.
D. Ask users to set their Google password to match their corporate password.

Answer: B

Explanation:
https://cloud.google.com/solutions/authenticating-corporate-users-in-a-hybrid-environment

NEW QUESTION 273
- (Topic 5)
You want to establish a Compute Engine application in a single VPC across two regions. The application must communicate over VPN to an on-premises network. How should you deploy the VPN?

A. Use VPC Network Peering between the VPC and the on-premises network.
B. Expose the VPC to the on-premises network using IAM and VPC Sharing.
C. Create a global Cloud VPN Gateway with VPN tunnels from each region to the on- premises peer gateway.
D. Deploy Cloud VPN Gateway in each regio
E. Ensure that each region has at least one VPN tunnel to the on-premises peer gateway.

Answer: C

Explanation:
https://cloud.google.com/vpn/docs/how-to/creating-static-vpns

NEW QUESTION 278
- (Topic 5)
Your company wants to try out the cloud with low risk. They want to archive approximately 100 TB of their log data to the cloud and test the analytics features available to them there, while also retaining that data as a long-term disaster recovery backup. Which two steps should they take? Choose 2 answers

A. Load logs into Google BigQuery.
B. Load logs into Google Cloud SQL.
C. Import logs into Google Stackdriver.
D. Insert logs into Google Cloud Bigtable.
E. Upload log files into Google Cloud Storage.

Answer: AE

NEW QUESTION 282
- (Topic 5)
Your customer wants to capture multiple GBs of aggregate real-time key performance indicators (KPIs) from their game servers running on Google Cloud Platform and monitor the KPIs with low latency. How should they capture the KPIs?

A. Store time-series data from the game servers in Google Bigtable, and view it using Google Data Studio.
B. Output custom metrics to Stackdriver from the game servers, and create a Dashboard in StackdriverMonitoring Console to view them.
C. Schedule BigQuery load jobs to ingest analytics files uploaded to Cloud Storage every ten minutes, and visualize the results in Google Data Studio.
D. Insert the KPIs into Cloud Datastore entities, and run ad hoc analysis and visualizations of them in Cloud Datalab.

Answer: A

Explanation:
https://cloud.google.com/monitoring/api/v3/metrics-details#metric-kinds

NEW QUESTION 285
- (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 288
- (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 289
- (Topic 5)
Your company has a project in Google Cloud with three Virtual Private Clouds (VPCs). There is a Compute Engine instance on each VPC. Network subnets do not overlap and must remain separated. The network configuration is shown below.

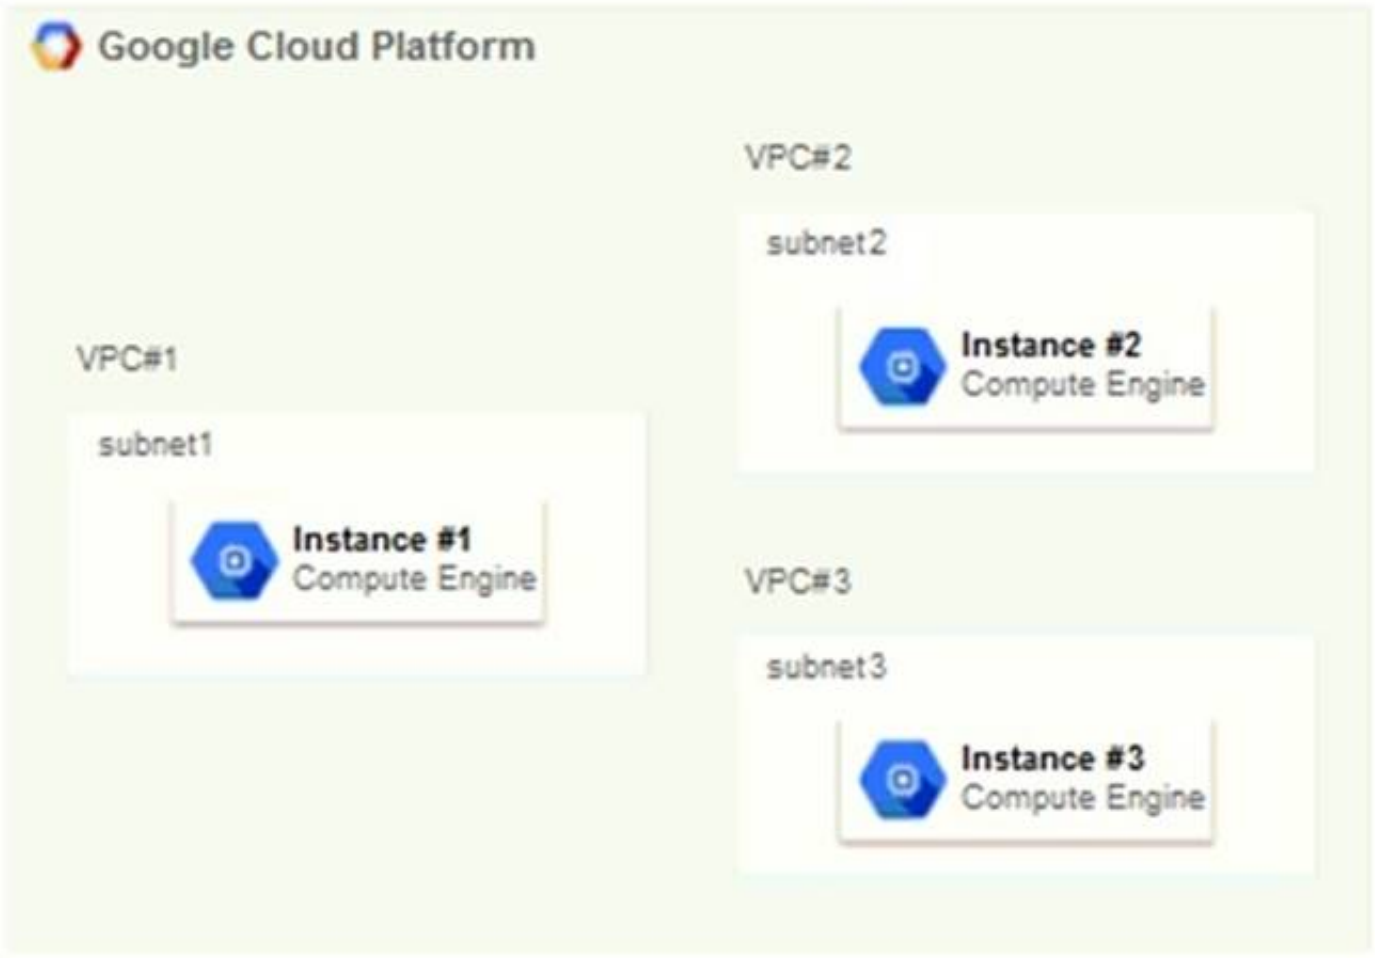

Instance #1 is an exception and must communicate directly with both Instance #2 and Instance #3 via internal IPs. How should you accomplish this?

A. Create a cloud router to advertise subnet #2 and subnet #3 to subnet #1.
B. Add two additional NICs to Instance #1 with the following configuration:•NIC1VPC: VPC #2SUBNETWORK: subnet #2•NIC2VPC: VPC #3SUBNETWORK: subnet #3Update firewall rules to enable traffic between instances.
C. Create two VPN tunnels via CloudVPN:•1 between VPC #1 and VPC #2.•1 between VPC #2 and VPC #3.Update firewall rules to enable traffic between the instances.
D. Peer all three VPCs:•Peer VPC #1 with VPC #2.•Peer VPC #2 with VPC #3.Update firewall rules to enable traffic between the instances.

Answer: B

Explanation:
As per GCP documentation: "By default, every instance in a VPC network has a single network interface. Use these instructions to create additional network interfaces. Each interface is attached to a different VPC network, giving that instance access to different VPC networks in Google Cloud. You cannot attach multiple network interfaces to the same VPC network." Refer to: https://cloud.google.com/vpc/docs/create-use-multiple-interfaces
https://cloud.google.com/vpc/docs/create-use-multiple- interfaces#i_am_not_able_to_connect_to_secondary_interfaces_internal_ip

NEW QUESTION 290
- (Topic 5)
Your company has developed a monolithic, 3-tier application to allow external users to upload and share files. The solution cannot be easily enhanced and lacks reliability. The development team would like to re-architect the application to adopt microservices and a fully managed service approach, but they need to convince their leadership that the effort is worthwhile. Which advantage(s) should they highlight to leadership?

A. The new approach will be significantly less costly, make it easier to manage the underlying infrastructure, and automatically manage the CI/CD pipelines.
B. The monolithic solution can be converted to a container with Docke
C. The generatedcontainer can then be deployed into a Kubernetes cluster.
D. The new approach will make it easier to decouple infrastructure from application, develop and release new features, manage the underlying infrastructure, manage CI/CD pipelines and perform A/B testing, and scale the solution if necessary.
E. The process can be automated with Migrate for Compute Engine.

Answer: C

Explanation:
The new approach will make it easier to decouple infrastructure from an application, develop and release new features, manage the underlying infrastructure, manage CI/CD pipelines and perform A/B testing, and scale the solution if necessary.

NEW QUESTION 294
- (Topic 5)
You have an application that makes HTTP requests to Cloud Storage. Occasionally the requests fail with

HTTP status codes of 5xx and 429.
How should you handle these types of errors?

A. Use gRPC instead of HTTP for better performance.
B. Implement retry logic using a truncated exponential backoff strategy.
C. Make sure the Cloud Storage bucket is multi-regional for geo-redundancy.
D. Monitor https://status.cloud.google.com/feed.atom and only make requests if Cloud Storage is not reportingan incident.

Answer: A

Explanation:
Reference https://cloud.google.com/storage/docs/json_api/v1/status-codes

NEW QUESTION 298
- (Topic 5)
Your company's user-feedback portal comprises a standard LAMP stack replicated across two zones. It is deployed in the us-central1 region and uses autoscaled managed instance groups on all layers, except the database. Currently, only a small group of select customers have access to the portal. The portal meets a 99.99% availability SLA under these conditions However next quarter, your company will be making the portal available to all users, including unauthenticated users. You need to develop a resiliency testing strategy to ensure the system maintains the SLA once they introduce additional user load. What should you do?

A. Capture existing users input, and replay captured user load until autoscale is triggered on all layer
B. At the same time, terminate all resources in one of the zones.
C. Create synthetic random user input, replay synthetic load until autoscale logic is triggered on at least one layer, and introduce "chaos" to the system by terminating random resources on both zones.
D. Expose the new system to a larger group of users, and increase group ' size each day until autoscale logic is tnggered on all layer
E. At the same time, terminate random resources on both zones.
F. Capture existing users input, and replay captured user load until resource utilization crosses 80%. Also, derive estimated number of users based on existing users usage of the app, and deploy enough resources to handle 200% of expected load.

Answer: A

NEW QUESTION 300
- (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 303
- (Topic 5)
Your company is using BigQuery as its enterprise data warehouse. Data is distributed over several Google Cloud projects. All queries on BigQuery need to be billed on a single
project. You want to make sure that no query costs are incurred on the projects that contain the data. Users should be able to query the datasets, but not edit them.
How should you configure users' access roles?

A. Add all users to a grou
B. Grant the group the role of BigQuery user on the billing project and BigQuerydataViewer on the projects that contain the data.
C. Add all users to a grou
D. Grant the group the roles of BigQuery dataViewer on the billing project andBigQuery user on the projects that contain the data.
E. Add all users to a grou
F. Grant the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data.
G. Add all users to a grou
H. Grant the group the roles of BigQuery dataViewer on the billing project andBigQuery jobUser on the projects that contain the data.

Answer: A

Explanation:
Reference: https://cloud.google.com/bigquery/docs/running-queries

NEW QUESTION 306
- (Topic 5)
Your applications will be writing their logs to BigQuery for analysis. Each application should have its own table.
Any logs older than 45 days should be removed. You want to optimize storage and follow Google recommended practices. What should you do?

A. Configure the expiration time for your tables at 45 days

B. Make the tables time-partitioned, and configure the partition expiration at 45 days
C. Rely on BigQuery's default behavior to prune application logs older than 45 days
D. Create a script that uses the BigQuery command line tool (bq) to remove records older than 45 days

Answer: B

Explanation:
https://cloud.google.com/bigquery/docs/managing-partitioned-tables

NEW QUESTION 308
- (Topic 5)
Your customer is moving their corporate applications to Google Cloud Platform. The security team wants detailed visibility of all projects in the organization. You provision the Google Cloud Resource Manager and set up yourself as the org admin. What Google Cloud Identity and Access Management (Cloud IAM) roles should you give to the security team'?

A. Org viewer, project owner
B. Org viewer, project viewer
C. Org admin, project browser
D. Project owner, network admin

Answer: B

Explanation:
https://cloud.google.com/iam/docs/using-iam-securely

NEW QUESTION 309
- (Topic 5)
Your customer is receiving reports that their recently updated Google App Engine application is taking approximately 30 seconds to load for some of their users. This behavior was not reported before the update. What strategy should you take?

A. Work with your ISP to diagnose the problem.
B. Open a support ticket to ask for network capture and flow data to diagnose the problem, then roll back your application.
C. Roll back to an earlier known good release initially, then use Stackdriver Trace and logging to diagnose the problem in a development/test/staging environment.
D. Roll back to an earlier known good release, then push the release again at a quieter period to investigat
E. Then use Stackdriver Trace and logging to diagnose the problem.

Answer: C

Explanation:
Stackdriver Logging allows you to store, search, analyze, monitor, and alert on log data and events from Google Cloud Platform and Amazon Web Services (AWS). Our API also allows ingestion of any custom log data from any source. Stackdriver Logging is a fully managed service that performs at scale and can ingest application and system log data from thousands of VMs. Even better, you can analyze all that log data in real time.
References: https://cloud.google.com/logging/

NEW QUESTION 314
- (Topic 5)
Your company has a Google Cloud project that uses BlgQuery for data warehousing There are some tables that contain personally identifiable information (PI!) Only the compliance team may access the PH. The other information in the tables must be available to the data science team. You want to minimize cost and the time it takes to assign appropriate access to the tables What should you do?

A. * 1 From the dataset where you have the source data, create views of tables that you want to share, excluding Pll* 2 Assign an appropriate project-level IAM role to the members of the data science team 3 Assign access controls to the dataset that contains the view
B. * 1 From the dataset where you have the source data, create materialized views of tables that you want to share excluding Pll* 2 Assign an appropriate project-level IAM role to the members of the data science team 3. Assign access controls to the dataset that contains the view.
C. * 1 Create a dataset for the data science team* 2 Create views of tables that you want to share excluding Pll* 3 Assign an appropriate project-level IAM role to the members of the data science team 4 Assign access controls to the dataset that contains the view* 5 Authorize the view to access the source dataset
D. * 1. Create a dataset for the data science team.* 2. Create materialized views of tables that you want to share, excluding Pll* 3. Assign an appropriate project-level IAM role to the members of the data science team * 4 Assign access controls to the dataset that contains the view* 5 Authorize the view to access the source dataset

Answer: C

Explanation:
This option can help minimize cost and time by using views and authorized datasets. Views are virtual tables defined by a SQL query that can exclude PII columns from the source tables. Views do not incur storage costs and do not duplicate data. Authorized datasets are datasets that have access to another dataset's data without granting direct access to individual users or groups. By creating a dataset for the data science team and creating views of tables that exclude PII, you can share only the relevant information with the team. By assigning an appropriate project-level IAM role to the members of the data science team, you can grant them access to the BigQuery service and resources. By assigning access controls to the dataset that contains the view, you can grant them access to query the views. By authorizing the view to access the source dataset, you can enable the view to read data from the source tables without exposing PII. The other options are not optimal for this scenario, because they either use materialized views instead of views, which incur storage costs and duplicate data (B, D), or do not create a separate dataset for the data science team, which makes it harder to manage access controls (A). References:
? https://cloud.google.com/bigquery/docs/views
? https://cloud.google.com/bigquery/docs/authorized-datasets

NEW QUESTION 319
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your Professional-Cloud-Architect Exam with Our Prep Materials Via below:

https://www.certleader.com/Professional-Cloud-Architect-dumps.html