perform an update to the application with minimal downtime to the application. What should you do?

A. Use kubect1 set image deployment/echo-deployment <new-image>
B. Use the rolling update functionality of the Instance Group behind the Kubernetes cluster
C. Update the deployment yaml file with the new container image. Use kubect1 delete deployment/echo-deployment and kubect1 create 瑄 <yaml-file>
D. Update the service yaml file which the new container image. Use kubect1 delete service/echoservice and kubect1 create 瑄 <yaml-file>

Correct Answer: A
Explanation:
https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps#updating_an_application

QUESTION 25
You are designing a Data Warehouse on Google Cloud and want to store sensitive data in BigQuery. Your company requires you to generate encryption keys outside of Google Cloud. You need to implement a solution. What should you do?

A. Generate a new key in Cloud Key Management Service (Cloud KMS). Store all data in Cloud Storage using the customer-managed key option and select the created key. Set up a Dataflow pipeline to decrypt the data and to store it in a BigQuery dataset.
B. Generate a new key in Cloud Key Management Service (Cloud KMS). Create a dataset in BigQuery using the customer-managed key option and select the created key
C. Import a key in Cloud KMS. Store all data in Cloud Storage using the customer-managed key option and select the created key. Set up a Dataflow pipeline to decrypt the data and to store it in a new BigQuery dataset.
D. Import a key in Cloud KMS. Create a dataset in BigQuery using the customer-supplied key option and select the created key.

Correct Answer: D
Explanation:
https://cloud.google.com/bigquery/docs/customer-managed-encryption

QUESTION 26
Your company plans to migrate a multi-petabyte data set to the cloud. The data set must be available 24hrs a day. Your business analysts have experience only with using a SQL interface. How should you store the data to optimize it for ease of analysis?

A. Load data into Google BigQuery.
B. Insert data into Google Cloud SQL.
C. Put flat files into Google Cloud Storage.
D. Stream data into Google Cloud Datastore.

Correct Answer: A
Explanation:
Google Big Query is for multi peta byte storage , HA(High availability) which means 24 hours, SQL interface.

https://medium.com/google-cloud/the-12-components-of-google-bigquery-c2b49829a7c7
https://cloud.google.com/solutions/bigquery-data-warehouse
https://cloud.google.com/bigquery/

BigQuery is Google's serverless, highly scalable, low cost enterprise data warehouse designed to make all your data analysts productive. Because there is no infrastructure to manage, you can focus on analyzing data to find meaningful insights using familiar SQL and you don't need a database administrator.

BigQuery enables you to analyze all your data by creating a logical data warehouse over managed, columnar storage as well as data from object storage, and spreadsheets.

https://cloud.google.com/bigquery/

QUESTION 27
You are working at a sports association whose members range in age from 8 to 30. The association collects a large amount of health data, such as sustained injuries. You are storing this data in BigQuery. Current legislation requires you to delete such information upon request of the subject. You want to design a solution that can accommodate such a request. What should you do?

A. Use a unique identifier for each individual. Upon a deletion request, delete all rows from BigQuery with this identifier.
B. When ingesting new data in BigQuery, run the data through the Data Loss Prevention (DLP) API to identify any personal information. As part of the DLP scan, save the result to Data Catalog. Upon a deletion request, query Data Catalog to find the column with personal information.
C. Create a BigQuery view over the table that contains all data. Upon a deletion request, exclude the rows that affect the subject's data from this view. Use this view instead of the source table for all analysis tasks.
D. Use a unique identifier for each individual. Upon a deletion request, overwrite the column with the unique identifier with a salted SHA256 of its value.

Correct Answer: B
Explanation:
Current legislation requires you to delete "SUCH" information upon request of the subject. " So from that point of view the question is not to delete the entire user records but specific data related to personal health data. With DLP you can use InfoTypes and InfoType detectors to specifically scan for those entries and how to act upon them (link https://cloud.google.com/dlp/docs/concepts-infotypes)

https://cloud.google.com/dlp#section-6

QUESTION 28
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.

Correct Answer: D

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

QUESTION 29
Your company has announced that they will be outsourcing operations functions. You want to allow developers to easily stage new versions of a cloud-based application in the production environment and allow the outsourced operations team to autonomously promote staged versions to production. You want to minimize the operational overhead of the solution. Which Google Cloud product should you migrate to?

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

Correct Answer: D
Explanation:
https://cloud.google.com/security/compliance/eba-outsourcing-mapping-gcp

QUESTION 30
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 country. Add analysts to their respective country-groups. Create a single group `all_analysts', and add all country-groups as members. Grant the `all-analysis' group the IAM role of BigQuery jobUser. Share the appropriate dataset with view access with each respective analyst country-group.
B.   Create a group per country. Add analysts to their respective country-groups. Create a single group `all_analysts', and add all country-groups as members. Grant the `all-analysis' group the IAM role of BigQuery jobUser. Share the appropriate tables with view access with each respective analyst countrygroup.
C.   Create a group per country. Add analysts to their respective country-groups. Create a single group `all_analysts', and add all country-groups as members. Grant the `all-analysis' group the IAM role of BigQuery dataViewer. Share the appropriate dataset with view access with each respective analyst country-group.
D.   Create a group per country. Add analysts to their respective country-groups. Create a single group `all_analysts', and add all country-groups as members. Grant the `all-analysis' group the IAM role of BigQuery dataViewer. Share the appropriate table with view access with each respective analyst countrygroup.

Correct Answer: A

QUESTION 31
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 GCP. Verify that direct peering can

establish a secure connection between your networks if Dedicated Interconnect fails.
B. Verify that Dedicated Interconnect can replicate files to GCP. Verify that Cloud VPN can establish a secure connection between your networks if Dedicated Interconnect fails.
C. Verify that the Transfer Appliance can replicate files to GCP. Verify that direct peering can establish a secure connection between your networks if the Transfer Appliance fails.
D. Verify that the Transfer Appliance can replicate files to GCP. Verify that Cloud VPN can establish a secure connection between your networks if the Transfer Appliance fails.

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

QUESTION 32
You want to allow your operations learn to store togs from all the production protects in your Organization, without during logs from other projects All of the production projects are contained in a folder. You want to ensure that all logs for existing and new production projects are captured automatically. What should you do?

A. Create an aggregated export on the Production folder. Set the log sink to be a Cloud Storage bucket in an operations project
B. Create an aggregated export on the Organization resource. Set the tog sink to be a Cloud Storage bucket in an operations project.
C. Create log exports in the production projects. Set the log sinks to be a Cloud Storage bucket in an operations project.
D. Create tog exports in the production projects. Set the tog sinks to be BigQuery datasets in the production projects and grant IAM access to the operations team to run queries on the datasets

Correct Answer: B

QUESTION 33
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.

Correct Answer: B
Explanation:
https://cloud.google.com/network-intelligence-center/docs/firewall-insights/how-to/using-firewall-insights

QUESTION 34
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 GDPR requirements.

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

QUESTION 35
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 Logging. Consult logs for Kubernetes Engine and Cloud SQL.
D. In the GCP Console, navigate to Cloud SQL. Restore the latest backup. Use kubect1 to restart all pods.

Correct Answer: C

QUESTION 36
You are using Cloud Shell and need to install a custom utility for use in a few weeks. Where can you store the file so it is in the default execution path and persists across sessions?

A. ~/bin
B. Cloud Storage
C. /google/scripts
D. /usr/local/bin

Correct Answer: D
Explanation:
https://medium.com/google-cloud/no-localhost-no-problem-using-google-cloud-shell-as-my-full-time-development-environment-22d5a1942439

QUESTION 37
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