

Google

Exam Questions Professional-Cloud-Architect

Google Certified Professional - Cloud Architect (GCP)



NEW QUESTION 1

- (Exam 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 2

- (Exam Topic 1)

For this question, refer to the Mountkirk Games case study.

Mountkirk Games' gaming servers are not automatically scaling properly. Last month, they rolled out a new feature, which suddenly became very popular. A record number of users are trying to use the service, but many of them are getting 503 errors and very slow response times. What should they investigate first?

- A. Verify that the database is online.
- B. Verify that the project quota hasn't been exceeded.
- C. Verify that the new feature code did not introduce any performance bugs.
- D. Verify that the load-testing team is not running their tool against production.

Answer: B

Explanation:

503 is service unavailable error. If the database was online everyone would get the 503 error. https://cloud.google.com/docs/quota#capping_usage

NEW QUESTION 3

- (Exam Topic 1)

For this question, refer to the Mountkirk Games case study.

Mountkirk Games has deployed their new backend on Google Cloud Platform (GCP). You want to create a thorough testing process for new versions of the backend before they are released to the public. You want the testing environment to scale in an economical way. How should you design the process?

- A. Create a scalable environment in GCP for simulating production load.
- B. Use the existing infrastructure to test the GCP-based backend at scale.
- C. Build stress tests into each component of your application using resources internal to GCP to simulate load.
- D. Create a set of static environments in GCP to test different levels of load — for example, high, medium, and low.

Answer: A

Explanation:

From scenario: Requirements for Game Backend Platform

- Dynamically scale up or down based on game activity
- Connect to a managed NoSQL database service
- Run customize Linux distro

NEW QUESTION 4

- (Exam 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 framework
- D. Focus on an API for the public.
- E. Use Google App Engine with the Swagger (open API Specification) framework
- F. Focus on an API for the public.
- G. Use Google Container Engine with a Django Python container
- H. Focus on an API for the public.
- I. Use Google Container Engine with a Tomcat container with the Swagger (Open API Specification) framework
- 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.-1712523

<https://cloud.google.com/endpoints/docs/openapi/architecture-overview> <https://cloud.google.com/storage/docs/gsutil/commands/test>

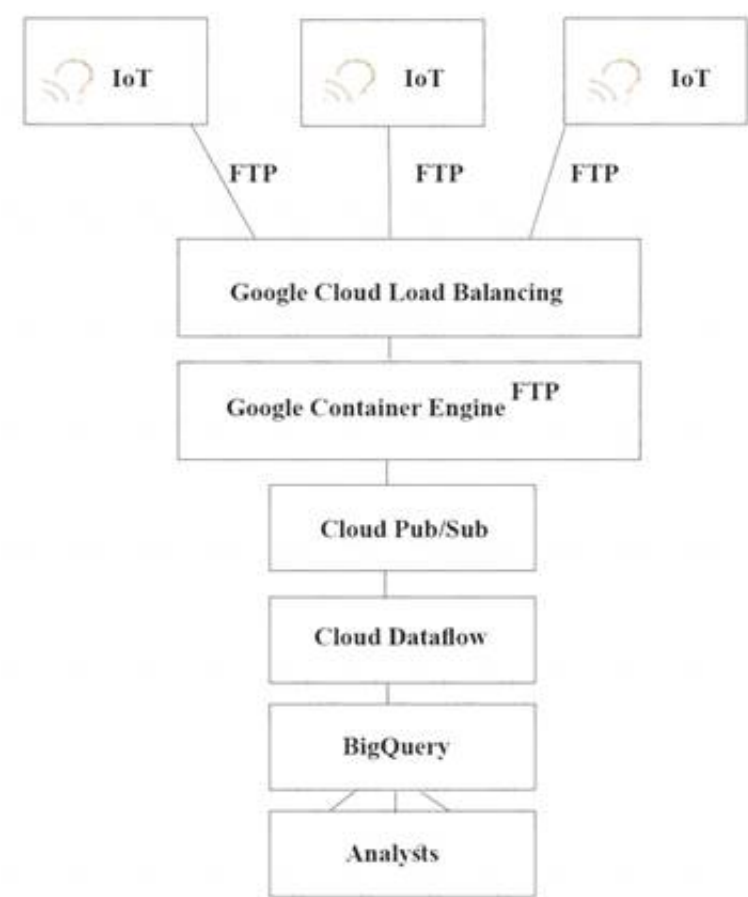
Develop, deploy, protect and monitor your APIs with Google Cloud Endpoints. Using an Open API Specification or one of our API frameworks, Cloud Endpoints gives you the tools you need for every phase of API development.

From scenario: Business Requirements

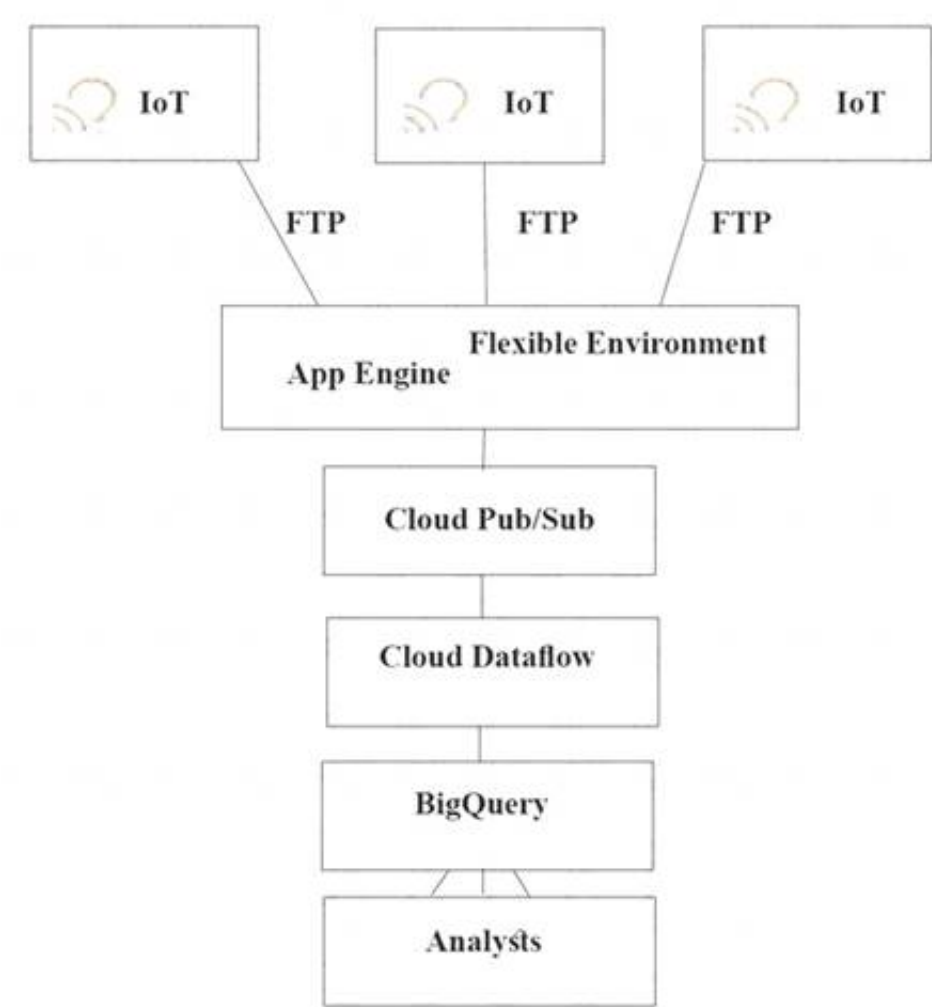
Decrease unplanned vehicle downtime to less than 1 week, without increasing the cost of carrying surplus inventory
Support the dealer network with more data on how their customers use their equipment to better position new products and services
Have the ability to partner with different companies – especially 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 5

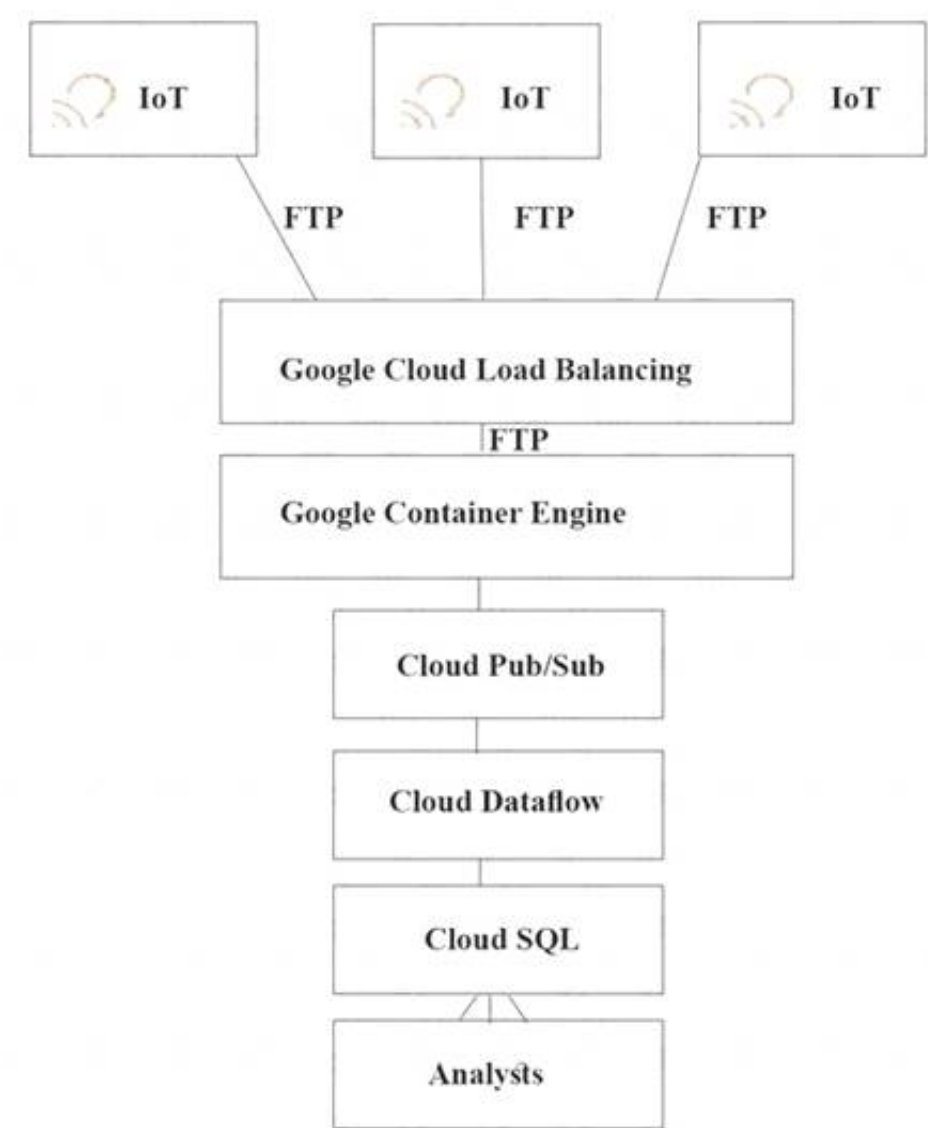
- (Exam Topic 2)
For this question, refer to the TerramEarth case study.
TerramEarth's CTO wants to use the raw data from connected vehicles to help identify approximately when a vehicle in the development team to focus their failure. You want to allow analysts to centrally query the vehicle data. Which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



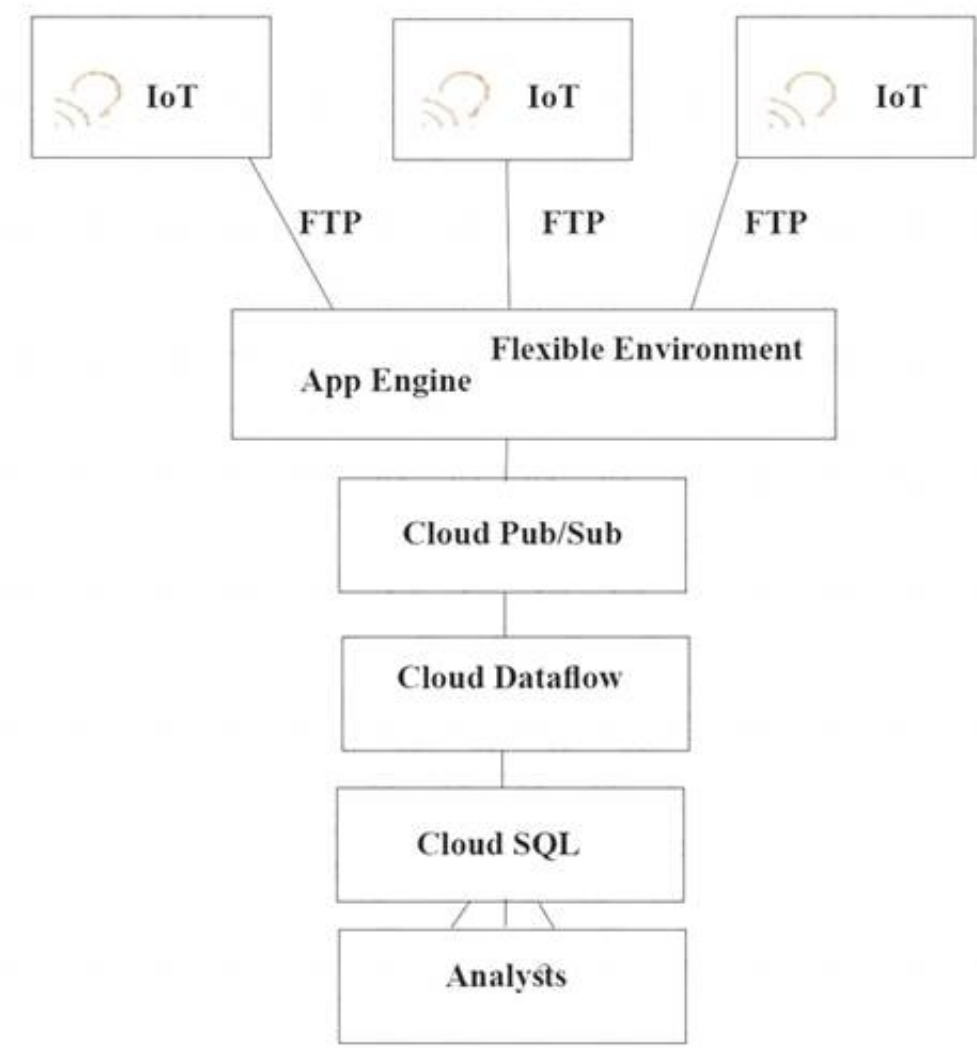
B)



C)



D)

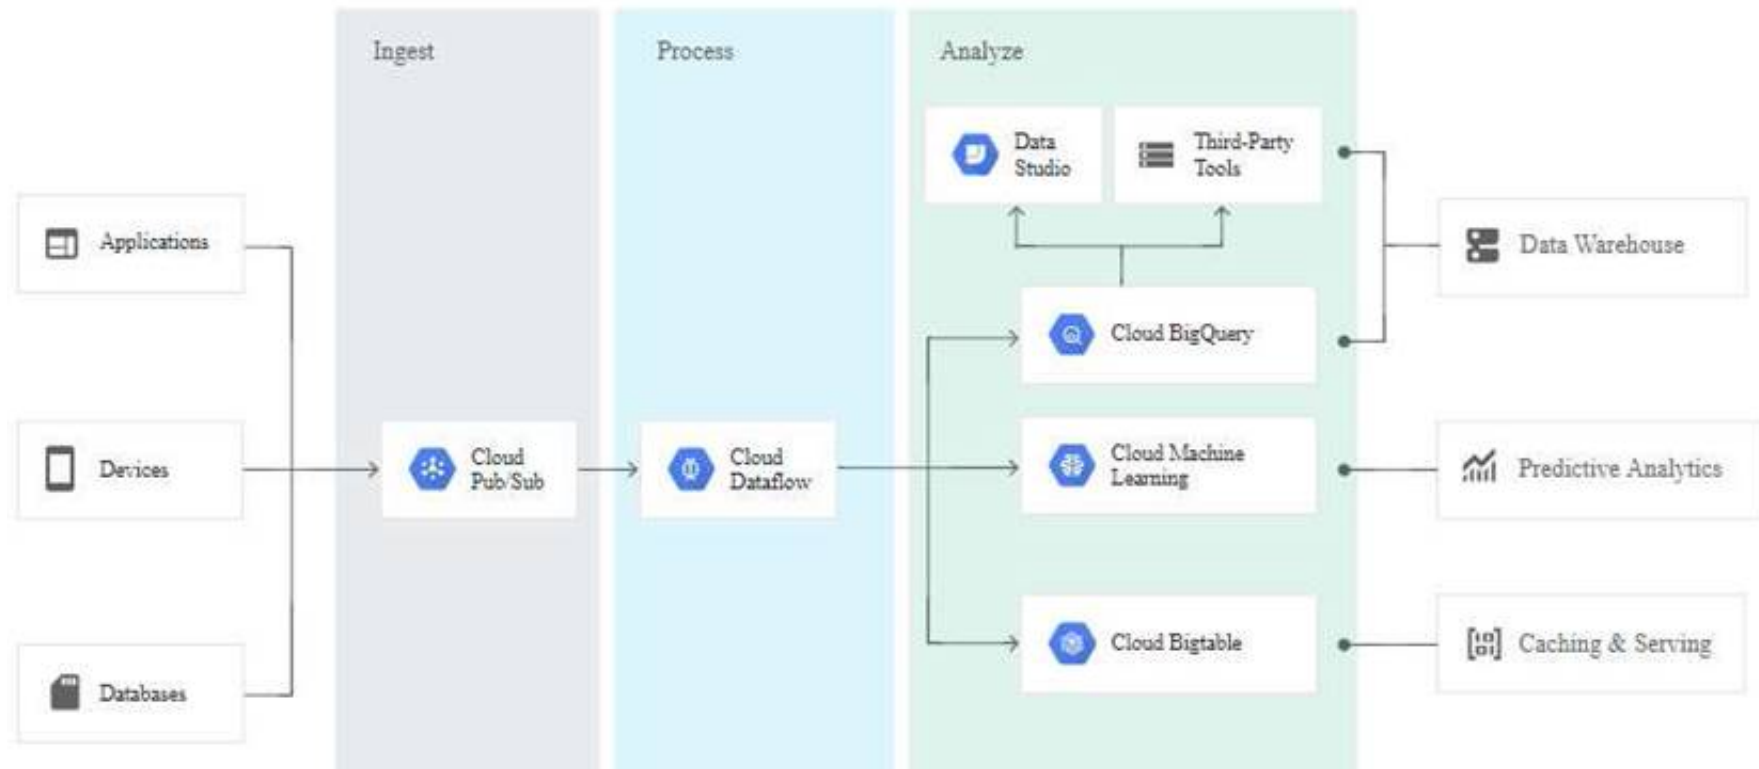


- A. Option A
- B. Option B
- C. Option C
- D. Option D

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>
<https://cloud.google.com/solutions/iot/>

NEW QUESTION 6

- (Exam Topic 2)

For this question, refer to the TerramEarth case study

Your development team has created a structured API to retrieve vehicle data. They want to allow third parties to develop tools for dealerships that use this vehicle event data. You want to support delegated authorization against this data. What should you do?

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

Answer: A

Explanation:

<https://cloud.google.com/appengine/docs/flexible/go/authorizing-apps> https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#delegate_application_autho Delegate application authorization with OAuth2

Cloud Platform APIs support OAuth 2.0, and scopes provide granular authorization over the methods that are supported. Cloud Platform supports both service-account and user-account OAuth, also called three-legged OAuth.

References:

https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#delegate_application_autho

<https://cloud.google.com/appengine/docs/flexible/go/authorizing-apps>

NEW QUESTION 7

- (Exam Topic 2)

For this question, refer to the TerramEarth case study.

To speed up data retrieval, more vehicles will be upgraded to cellular connections and be able to transmit data to the ETL process. The current FTP process is error-prone and restarts the data transfer from the start of the file when connections fail, which happens often. You want to improve the reliability of the solution and minimize data transfer time on the cellular connections. What should you do?

- A. Use one Google Container Engine cluster of FTP server
- B. Save the data to a Multi-Regional bucket
- C. Run the ETL process using data in the bucket.
- D. Use multiple Google Container Engine clusters running FTP servers located in different region
- E. Save the data to Multi-Regional buckets in us, eu, and asi
- F. Run the ETL process using the data in the bucket.
- G. Directly transfer the files to different Google Cloud Multi-Regional Storage bucket locations in us, eu, and asia using Google APIs over HTTP(S). Run the ETL process using the data in the bucket.
- H. Directly transfer the files to a different Google Cloud Regional Storage bucket location in us, eu, and asia using Google APIs over HTTP(S). Run the ETL process to retrieve the data from each Regional bucket.

Answer: D

Explanation:

<https://cloud.google.com/storage/docs/locations>

NEW QUESTION 8

- (Exam Topic 2)

For this question refer to the TerramEarth case study

Operational parameters such as oil pressure are adjustable on each of TerramEarth's vehicles to increase their efficiency, depending on their environmental conditions. Your primary goal is to increase the operating efficiency of all 20 million cellular and unconnected vehicles in the field How can you accomplish this goal?

- A. Have your engineers inspect the data for patterns, and then create an algorithm with rules that make operational adjustments automatically.
- B. Capture all operating data, train machine learning models that identify ideal operations, and run locally to make operational adjustments automatically.
- C. Implement a Google Cloud Dataflow streaming job with a sliding window, and use Google Cloud Messaging (GCM) to make operational adjustments automatically.
- D. Capture all operating data, train machine learning models that identify ideal operations, and host in Google Cloud Machine Learning (ML) Platform to make operational adjustments automatically.

Answer: B

NEW QUESTION 9

- (Exam 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 10

- (Exam Topic 3)

For this question, refer to the JencoMart case study.

JencoMart wants to move their User Profiles database to Google Cloud Platform. Which Google Database should they use?

- A. Cloud Spanner
- B. Google BigQuery
- C. Google Cloud SQL
- D. Google Cloud Datastore

Answer: D

Explanation:

<https://cloud.google.com/datastore/docs/concepts/overview> Common workloads for Google Cloud Datastore:

- User profiles
- Product catalogs
- Game state

References: <https://cloud.google.com/storage-options/> <https://cloud.google.com/datastore/docs/concepts/overview>

NEW QUESTION 10

- (Exam 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 and one for production. Each account should contain one project for every application.
- B. 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.
- C. Create a single G Suite account to manage users with each stage of each application in its own project.
- D. 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 14

- (Exam 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, Dress4win wants to distribute its 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 the regional 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 of virtual machines managed by your operations team.
- C. Use regional managed instance groups and a global load balancer to increase reliability by providing automatic 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 of virtual machines as part of a separate managed instance groups.

Answer: A

NEW QUESTION 18

- (Exam 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 19

- (Exam Topic 4)

For this question, refer to the Dress4Win case study.

You want to ensure Dress4Win's sales and tax records remain available for infrequent viewing by auditors for at least 10 years. Cost optimization is your top priority. Which cloud services should you choose?

- A. Google Cloud Storage Coldline to store the data, and gsutil to access the data.
- B. Google Cloud Storage Nearline to store the data, and gsutil to access the data.
- C. Google Bigtable with US or EU as location to store the data, and gcloud to access the data.
- D. BigQuery to store the data, and a web server cluster in a managed instance group to access the data. Google Cloud SQL mirrored across two distinct regions to store the data, and a Redis cluster in a managed instance group to access the data.

Answer: A

Explanation:

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

NEW QUESTION 22

- (Exam Topic 4)

Dress4win has end to end tests covering 100% of their endpoints.

They want to ensure that the move of cloud does not introduce any new bugs.

Which additional testing methods should the developers employ to prevent an outage?

- A. They should run the end to end tests in the cloud staging environment to determine if the code is working as intended.
- B. They should enable google stack driver debugger on the application code to show errors in the code
- C. They should add additional unit tests and production scale load tests on their cloud staging environment.
- 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 27

- (Exam 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 31

- (Exam Topic 4)

For this question, refer to the Dress4Win case study.

As part of their new application experience, Dress4Wm allows customers to upload images of themselves. The customer has exclusive control over who may view these images. Customers should be able to upload images with minimal latency and also be shown their images quickly on the main application page when they log in. Which configuration should Dress4Win use?

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

- I. As storage needs increase, add more persistent disks and/or node
- J. Use a Google Cloud SQL database to maintain metadata that maps each customer's ID to their image files.

Answer: A

NEW QUESTION 34

- (Exam Topic 4)

For this question, refer to the Dress4Win case study.

Dress4Win would like to become familiar with deploying applications to the cloud by successfully deploying some applications quickly, as is. They have asked for your recommendation. What should you advise?

- A. Identify self-contained applications with external dependencies as a first move to the cloud.
- B. Identify enterprise applications with internal dependencies and recommend these as a first move to the cloud.
- C. Suggest moving their in-house databases to the cloud and continue serving requests to on-premise applications.
- D. Recommend moving their message queuing servers to the cloud and continue handling requests to on-premise applications.

Answer: A

Explanation:

<https://cloud.google.com/blog/products/gcp/the-five-phases-of-migrating-to-google-cloud-platform>

NEW QUESTION 36

- (Exam Topic 5)

As part of implementing their disaster recovery plan, your company is trying to replicate their production MySQL database from their private data center to their GCP project using a Google Cloud VPN connection. They are experiencing latency issues and a small amount of packet loss that is disrupting the replication.

What should they do?

- A. Configure their replication to use UDP.
- B. Configure a Google Cloud Dedicated Interconnect.
- C. Restore their database daily using Google Cloud SQL.
- D. Add additional VPN connections and load balance them.
- E. Send the replicated transaction to Google Cloud Pub/Sub.

Answer: B

NEW QUESTION 39

- (Exam Topic 5)

You have an application deployed on Kubernetes Engine using a Deployment named echo-deployment. The deployment is exposed using a Service called echo-service. You need to perform an update to the application with minimal downtime to the application. What should you do?

- A. Use kubectl 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
- D. Use kubectl delete deployment/ echo-deployment and kubectl create -f <yaml-file>
- E. Update the service yaml file with the new container image
- F. Use kubectl delete service/echoservice and kubectl create -f <yaml-file>

Answer: A

Explanation:

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

NEW QUESTION 40

- (Exam 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 43

- (Exam Topic 5)

You have developed an application using Cloud ML Engine that recognizes famous paintings from uploaded images. You want to test the application and allow specific people to upload images for the next 24 hours. Not all users have a Google Account. How should you have users upload images?

- A. Have users upload the images to Cloud Storage
- B. Protect the bucket with a password that expires after 24 hours.
- C. Have users upload the images to Cloud Storage using a signed URL that expires after 24 hours.
- D. Create an App Engine web application where users can upload image

- E. Configure App Engine to disable the application after 24 hour
- F. Authenticate users via Cloud Identity.
- G. Create an App Engine web application where users can upload images for the next 24 hours. Authenticate users via Cloud Identity.

Answer: A

Explanation:

<https://cloud.google.com/blog/products/storage-data-transfer/uploading-images-directly-to-cloud-storage-by-usi>

NEW QUESTION 47

- (Exam 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 Stackdriver Monitoring 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 48

- (Exam Topic 5)

You are using a single Cloud SQL instance to serve your application from a specific zone. You want to introduce high availability. What should you do?

- A. Create a read replica instance in a different region
- B. Create a failover replica instance in a different region
- C. Create a read replica instance in the same region, but in a different zone
- D. Create a failover replica instance in the same region, but in a different zone

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/mysql/high-availability>

NEW QUESTION 51

- (Exam Topic 5)

You set up an autoscaling instance group to serve web traffic for an upcoming launch. After configuring the instance group as a backend service to an HTTP(S) load balancer, you notice that virtual machine (VM) instances are being terminated and re-launched every minute. The instances do not have a public IP address. You have verified the appropriate web response is coming from each instance using the curl command. You want to ensure the backend is configured correctly. What should you do?

- A. Ensure that a firewall rule exists to allow source traffic on HTTP/HTTPS to reach the load balancer.
- B. Assign a public IP to each instance and configure a firewall rule to allow the load balancer to reach the instance public IP.
- C. Ensure that a firewall rule exists to allow load balancer health checks to reach the instances in the instance group.
- D. Create a tag on each instance with the name of the load balance
- E. Configure a firewall rule with the name of the load balancer as the source and the instance tag as the destination.

Answer: C

Explanation:

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

The best practice when configuration a health check is to check health and serve traffic on the same port. However, it is possible to perform health checks on one port, but serve traffic on another. If you do use two different ports, ensure that firewall rules and services running on instances are configured appropriately. If you run health checks and serve traffic on the same port, but decide to switch ports at some point, be sure to update both the backend service and the health check. Backend services that do not have a valid global forwarding rule referencing it will not be health checked and will have no health status.

References: <https://cloud.google.com/compute/docs/load-balancing/http/backend-service>

NEW QUESTION 52

- (Exam Topic 5)

One of the developers on your team deployed their application in Google Container Engine with the Dockerfile below. They report that their application deployments are taking too long.

```
FROM ubuntu:16.04

COPY . /src

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

RUN pip install -r requirements.txt
```

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

- A. Remove Python after running pip.
- B. Remove dependencies from requirements.txt.
- C. Use a slimmed-down base image like Alpine linux.

- D. Use larger machine types for your Google Container Engine node pools.
- E. Copy the source after the package dependencies (Python and pip) are installed.

Answer: CE

Explanation:

The speed of deployment can be changed by limiting the size of the uploaded app, limiting the complexity of the build necessary in the Dockerfile, if present, and by ensuring a fast and reliable internet connection.

Note: Alpine Linux is built around musl libc and busybox. This makes it smaller and more resource efficient than traditional GNU/Linux distributions. A container requires no more than 8 MB and a minimal installation to disk requires around 130 MB of storage. Not only do you get a fully-fledged Linux environment but a large selection of packages from the repository.

References: <https://groups.google.com/forum/#!topic/google-appengine/hZMEkmmObDU> <https://www.alpinelinux.org/about/>

NEW QUESTION 53

- (Exam 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 explore and clean your data.

Answer: B

Explanation:

<https://cloud.google.com/dataprep/>

NEW QUESTION 54

- (Exam Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 55

- (Exam Topic 5)

Your company is designing its data lake on Google Cloud and wants to develop different ingestion pipelines to collect unstructured data from different sources. After the data is stored in Google Cloud, it will be processed in several data pipelines to build a recommendation engine for end users on the website. The structure of the data retrieved from the source systems can change at any time. The data must be stored exactly as it was retrieved for reprocessing purposes in case the data structure is incompatible with the current processing pipelines. You need to design an architecture to support the use case after you retrieve the data. What should you do?

- A. Send the data through the processing pipeline, and then store the processed data in a BigQuery table for reprocessing.
- B. Store the data in a BigQuery tabl
- C. Design the processing pipelines to retrieve the data from the table.
- D. Send the data through the processing pipeline, and then store the processed data in a Cloud Storage bucket for reprocessing.
- E. Store the data in a Cloud Storage bucke
- F. Design the processing pipelines to retrieve the data from the bucket

Answer: D

NEW QUESTION 59

- (Exam Topic 5)

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

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

- E. Configure Private Google Access on the Virtual Private Cloud (VPC)
- F. Configure the GKE cluster as a private cluster
- G. Configure Private Google Access on the Virtual Private Cloud (VPC)

Answer: B

Explanation:

A Cloud NAT gateway can perform NAT for nodes and Pods in a private cluster, which is a type of VPC-native cluster. The Cloud NAT gateway must be configured to apply to at least the following subnet IP address ranges for the subnet that your cluster uses:

- Subnet primary IP address range (used by nodes)
- Subnet secondary IP address range used for Pods in the cluster
- Subnet secondary IP address range used for Services in the cluster

The simplest way to provide NAT for an entire private cluster is to configure a Cloud NAT gateway to apply to all of the cluster's subnet's IP address ranges.

<https://cloud.google.com/nat/docs/overview>

NEW QUESTION 64

- (Exam 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 69

- (Exam 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 data
- B. Lifecycle-managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.
- C. Memcache backed by Cloud Datastore for the customer session state data
- D. Lifecycle-managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.
- E. Memcache backed by Cloud SQL for customer session state data
- 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 data
- I. Assorted local SSD-backed 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 70

- (Exam Topic 5)

Your company is moving 75 TB of data into Google Cloud. You want to use Cloud Storage and follow Google-recommended practices. What should you do?

- A. Move your data onto a Transfer Appliance
- B. Use a Transfer Appliance Rehydrator to decrypt the data into Cloud Storage.
- C. Move your data onto a Transfer Appliance
- D. Use Cloud Dataprep to decrypt the data into Cloud Storage.
- E. Install gsutil on each server that contains data
- F. Use resumable transfers to upload the data into Cloud Storage.
- G. Install gsutil on each server containing data
- H. Use streaming transfers to upload the data into Cloud Storage.

Answer: A

Explanation:

<https://cloud.google.com/transfer-appliance/docs/2.0/faq>

NEW QUESTION 73

- (Exam Topic 5)

Auditors visit your teams every 12 months and ask to review all the Google Cloud Identity and Access Management (Cloud IAM) policy changes in the previous 12 months. You want to streamline and expedite the analysis and audit process. What should you do?

- A. Create custom Google Stackdriver alerts and send them to the auditor.
- B. Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor.

- C. Use cloud functions to transfer log entries to Google Cloud SQL and use ACLS and views to limit an auditor's view.
- D. Enable Google Cloud Storage (GCS) log export to audit logs into a GCS bucket and delegate access to the bucket.

Answer: D

Explanation:

Export the logs to Google Cloud Storage bucket - Archive Storage, as it will not be used for 1 year, price for which is \$0.004 per GB per Month. The price for long term storage in BigQuery is \$0.01 per GB per Month (250% more). Also for analysis purpose, whenever Auditors are there (once per year), you can use BigQuery and use GCS bucket as external data source. BigQuery supports querying Cloud Storage data from these storage classes:
Standard Nearline Coldline Archive

NEW QUESTION 75

- (Exam 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 archived files and files uploaded daily using the gsutil -m option.
- B. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transfer archived data to Cloud Storage.
- C. Establish a connection with Google using a Dedicated Interconnect or Direct 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 transfer archived data to Cloud Storage.
- E. Establish one Cloud VPN Tunnel to VPC networks over the public internet, and compress 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 Storage.
- 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 80

- (Exam 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 cluster.
- B. Use Deployment Manager to create the deployment.
- C. Use gcloud to create a Kubernetes cluster.
- D. Use kubectl to create the deployment.
- E. Use kubectl to create a Kubernetes cluster.
- F. Use Deployment Manager to create the deployment.
- G. Use kubectl to create a Kubernetes cluster.
- H. Use kubectl to create the deployment.

Answer: B

Explanation:

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

NEW QUESTION 81

- (Exam Topic 5)

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

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

Answer: C

NEW QUESTION 82

- (Exam Topic 5)

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.
- B. Set up a Dataflow pipeline to decrypt the data and to store it in a BigQuery dataset.
- C. 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.
- D. Import a key in Cloud KMS.
- E. Store all data in Cloud Storage using the customer-managed key option and select the created key.
- F. Set up a Dataflow pipeline to decrypt the data and to store it in a new BigQuery dataset.
- G. Import a key in Cloud KMS.
- H. Create a dataset in BigQuery using the customer-supplied key option and select the created key.

Answer: D

Explanation:

<https://cloud.google.com/bigquery/docs/customer-managed-encryption>

NEW QUESTION 83

- (Exam 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 86

- (Exam 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 90

- (Exam Topic 5)

You are implementing the infrastructure for a web service on Google Cloud. The web service needs to receive and store the data from 500,000 requests per second. The data will be queried later in real time, based on exact matches of a known set of attributes. There will be periods where the web service will not receive any requests. The business wants to keep costs low. Which web service platform and database should you use for the application?

- A. Cloud Run and BigQuery
- B. Cloud Run and Cloud Bigtable
- C. A Compute Engine autoscaling managed instance group and BigQuery
- D. A Compute Engine autoscaling managed instance group and Cloud Bigtable

Answer: B

Explanation:

<https://cloud.google.com/run/docs/about-instance-autoscaling> <https://cloud.google.com/blog/topics/developers-practitioners/bigtable-vs-bigquery-whats-difference>

NEW QUESTION 91

- (Exam Topic 5)

Your company just finished a rapid lift and shift to Google Compute Engine for your compute needs. You have another 9 months to design and deploy a more cloud-native solution. Specifically, you want a system that is no-ops and auto-scaling. Which two compute products should you choose? Choose 2 answers

- A. Compute Engine with containers
- B. Google Kubernetes Engine with containers
- C. Google App Engine Standard Environment
- D. Compute Engine with custom instance types
- E. Compute Engine with managed instance groups

Answer: BC

Explanation:

B: With Container Engine, Google will automatically deploy your cluster for you, update, patch, secure the nodes.

Kubernetes Engine's cluster autoscaler automatically resizes clusters based on the demands of the workloads you want to run.

C: Solutions like Datastore, BigQuery, AppEngine, etc are truly NoOps.

App Engine by default scales the number of instances running up and down to match the load, thus providing consistent performance for your app at all times while minimizing idle instances and thus reducing cost.

Note: At a high level, NoOps means that there is no infrastructure to build out and manage during usage of the platform. Typically, the compromise you make with NoOps is that you lose control of the underlying infrastructure.

References:

<https://www.quora.com/How-well-does-Google-Container-Engine-support-Google-Cloud-Platform%E2%80%99s>

NEW QUESTION 94

- (Exam 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 group
- D. Add his SSH key to the project-wide SSH keys
- E. Disable autoscaling for the instance group
- 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 99

- (Exam 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 104

- (Exam Topic 5)

Your architecture calls for the centralized collection of all admin activity and VM system logs within your project. How should you collect these logs from both VMs and services?

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

Answer: B

Explanation:

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

NEW QUESTION 105

- (Exam 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 training
- B. Set a deadline for the new GCP project.
- C. Allocate budget for team training
- 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 107

- (Exam 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 112

- (Exam Topic 5)

You are analyzing and defining business processes to support your startup's trial usage of GCP, and you don't yet know what consumer demand for your product will be. Your manager requires you to minimize GCP service costs and adhere to Google best practices. What should you do?

- A. Utilize free tier and sustained use discount
- B. Provision a staff position for service cost management.
- C. Utilize free tier and sustained use discount
- D. Provide training to the team about service cost management.
- E. Utilize free tier and committed use discount
- F. Provision a staff position for service cost management.
- G. Utilize free tier and committed use discount
- H. Provide training to the team about service cost management.

Answer: D

Explanation:

https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#billing_and_management

NEW QUESTION 116

- (Exam 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 bucket
- 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 bucket
- 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 117

- (Exam 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 branch
- 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 branch
- 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 Registry
- F. Create a deployment pipeline that watches for new images and deploys the new image on the development cluster
- 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 Registry
- I. Rely on Vulnerability Scanning to ensure the code tests succeed
- J. As the final step of the Cloud Build process, deploy the new container image on the development cluster
- 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 120

- (Exam Topic 5)

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

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

Answer: AC

NEW QUESTION 124

- (Exam Topic 5)

You are designing a large distributed application with 30 microservices. Each of your distributed microservices needs to connect to a database back-end. You want to store the credentials securely. Where should you store the credentials?

- A. In the source code
- B. In an environment variable
- C. In a secret management system
- D. In a config file that has restricted access through ACLs

Answer: C

Explanation:

https://cloud.google.com/docs/authentication/production#providing_credentials_to_your_application

NEW QUESTION 127

- (Exam Topic 5)

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

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

Answer: D

Explanation:

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

NEW QUESTION 130

- (Exam Topic 5)

You write a Python script to connect to Google BigQuery from a Google Compute Engine virtual machine. The script is printing errors that it cannot connect to BigQuery. What should you do to fix the script?

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

Answer: B

Explanation:

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

<https://cloud.google.com/compute/docs/access/service-accounts>

NEW QUESTION 134

- (Exam Topic 5)

Your company has an application running on Google Cloud that is collecting data from thousands of physical devices that are globally distributed. Data is publish to Pub/Sub and streamed in real time into an SSO Cloud Bigtable cluster via a Dataflow pipeline. The operations team informs you that your Cloud Bigtable cluster has a hot-spot and queries are taking longer man expected You need to resolve the problem and prevent it from happening in the future What should you do?

- A. Advise your clients to use HBase APIs instead of NodeJS APIs.
- B. Review your RowKey strategy and ensure that keys are evenly spread across the alphabet.
- C. Delete records older than 30 days.
- D. Double the number of nodes you currently have.

Answer: B

NEW QUESTION 135

- (Exam Topic 5)

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

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

Answer: A

NEW QUESTION 137

- (Exam 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 139

- (Exam 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 latency
- 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 cluster
- 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 142

- (Exam 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 triggered 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 144

- (Exam 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 reporting an incident.

Answer: A

Explanation:

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

NEW QUESTION 145

- (Exam 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 149

- (Exam Topic 5)

Your company has an application running on Compute Engine that 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 load balancer and configure it with two backends* 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 music 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 files directly from the backend Compute Engine instance

Answer: B

NEW QUESTION 151

- (Exam Topic 5)

Your operations team currently stores 10 TB of data in 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. Which approach should they use?

- A. Use the `gsutil mv` command to 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-optio> <https://cloud.google.com/storage-transfer-service>

NEW QUESTION 152

- (Exam Topic 5)

You are deploying a PHP App Engine Standard service with SQL as the backend. You want to minimize the number of queries to the database. What should you do?

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

Answer: A

Explanation:

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

NEW QUESTION 153

- (Exam 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 group
- B. Grant the group the role of BigQuery user on the billing project and BigQuery dataViewer on the projects that contain the data.
- C. Add all users to a group

- D. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery user on the projects that contain the data.
- E. Add all users to a group
- 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 group
- H. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery jobUser on the projects that contain the data.

Answer: A

Explanation:

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

NEW QUESTION 154

- (Exam Topic 5)

Your company has an application running on multiple Compute Engine instances. You need to ensure that the application can communicate with an on-premises service that requires high throughput via internal IPs, while minimizing latency. What should you do?

- A. Use OpenVPN to configure a VPN tunnel between the on-premises environment and Google Cloud.
- B. Configure a direct peering connection between the on-premises environment and Google Cloud.
- C. Use Cloud VPN to configure a VPN tunnel between the on-premises environment and Google Cloud.
- D. Configure a Cloud Dedicated Interconnect connection between the on-premises environment and Google Cloud.

Answer: C

Explanation:

Reference <https://cloud.google.com/architecture/setting-up-private-access-to-cloud-apis-through-vpn-tunnels>

NEW QUESTION 156

- (Exam 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 instances
- 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 instances
- 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 traffic
- F. Rely on the implied deny egress rule with priority 100 to block all traffic for all instances.
- G. Create an egress rule with priority 100 to allow the Active Directory traffic
- 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 158

- (Exam 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 161

- (Exam 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 years
- 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 bucket
- 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 Writer. Use the service account to upload new files.

Answer: A

Explanation:

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

NEW QUESTION 165

- (Exam 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 168

- (Exam 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
- E. * 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 172

- (Exam 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 176

- (Exam 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 ls -lr gs://backups/** to find and remove items older than 90 days.
- D. Schedule a cron script using gsutil ls -l 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 177

- (Exam Topic 5)

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

- A. App Engine
- B. Cloud Endpoints

- C. Compute Engine
- D. Google Kubernetes Engine

Answer: A

Explanation:

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

NEW QUESTION 180

- (Exam 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 region. 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.
- B. Deploy the application on a Compute Engine instance
- C. 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.
- D. Deploy the application on two Compute Engine instance groups, each in the same project but in a different region
- E. 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.
- F. Deploy the application on two Compute Engine instance groups, each in separate project and a different region
- G. Use the first instance group to serve traffic, and use the HTTP load balancing service to fail over to the standby instance in case of a disaster.

Answer: C

NEW QUESTION 184

- (Exam Topic 5)

The application reliability team at your company has added a debug feature to their backend service to send all server events to Google Cloud Storage for eventual analysis. The event records are at least 50 KB and at most 15 MB and are expected to peak at 3,000 events per second. You want to minimize data loss. Which process should you implement?

- A. • Append metadata to file body. • Compress individual files. • Name files with serverName-Timestamp. • Create a new bucket if bucket is older than 1 hour and save individual files to the new bucket
- B. Otherwise, save files to existing bucket
- C. • Batch every 10,000 events with a single manifest file for metadata. • Compress event files and manifest file into a single archive file. • Name files using serverName-EventSequence. • Create a new bucket if bucket is older than 1 day and save the single archive file to the new bucket
- D. Otherwise, save the single archive file to existing bucket.
- E. • Compress individual files. • Name files with serverName-EventSequence. • Save files to one bucket • Set custom metadata headers for each object after saving.
- F. • Append metadata to file body. • Compress individual files. • Name files with a random prefix pattern. • Save files to one bucket

Answer: D

Explanation:

In order to maintain a high request rate, avoid using sequential names. Using completely random object names will give you the best load distribution. Randomness after a common prefix is effective under the prefix <https://cloud.google.com/storage/docs/request-rate>

NEW QUESTION 185

- (Exam 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 186

- (Exam Topic 5)

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

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

Pods to automatically deploy new production releases.

Answer: D

Explanation:

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

NEW QUESTION 191

- (Exam Topic 5)

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

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

Answer: D

Explanation:

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

NEW QUESTION 196

- (Exam 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_orga

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

NEW QUESTION 198

- (Exam Topic 5)

You want to create a private connection between your instances on Compute Engine and your on-premises data center. You require a connection of at least 20 Gbps. You want to follow Google-recommended practices.

How should you set up the connection?

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

Answer: A

Explanation:

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

NEW QUESTION 201

- (Exam Topic 5)

Your company has a support ticketing solution that uses App Engine Standard. The project that contains the App Engine application already has a Virtual Private Cloud(VPC) network fully connected to the company's on-premises environment through a Cloud VPN tunnel. You want to enable App Engine application to communicate with a database that is running in the company's on-premises environment. What should you do?

- A. Configure private services access
- B. Configure private Google access for on-premises hosts only
- C. Configure serverless VPC access
- D. Configure private Google access

Answer: C

Explanation:

<https://cloud.google.com/appengine/docs/standard/python3/connecting-vpc> https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databases#on_premises

NEW QUESTION 202

- (Exam Topic 5)

Your customer wants to do resilience testing of their authentication layer. This consists of a regional managed instance group serving a public REST API that reads from and writes to a Cloud SQL instance.

What should you do?

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

Answer: C

NEW QUESTION 204

- (Exam Topic 5)

You want your Google Kubernetes Engine cluster to automatically add or remove nodes based on CPU load. What should you do?

- A. Configure a HorizontalPodAutoscaler with a target CPU usage
- B. Enable the Cluster Autoscaler from the GCP Console.
- C. Configure a HorizontalPodAutoscaler with a target CPU usage
- D. Enable autoscaling on the managed instance group for the cluster using the gcloud command.
- E. Create a deployment and set the maxUnavailable and maxSurge properties
- F. Enable the Cluster Autoscaler using the gcloud command.
- G. Create a deployment and set the maxUnavailable and maxSurge properties
- H. Enable autoscaling on the cluster managed instance group from the GCP Console.

Answer: B

NEW QUESTION 207

- (Exam 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 GCP
- B. Verify that direct peering can establish a secure connection between your networks if Dedicated Interconnect fails.
- C. Verify that Dedicated Interconnect can replicate files to GCP
- 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 GCP
- F. Verify that direct peering can establish a secure connection between your networks if the Transfer Appliance fails.
- G. Verify that the Transfer Appliance can replicate files to GCP
- 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

- (Exam Topic 5)

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

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

Answer: B

NEW QUESTION 215

- (Exam Topic 5)

You are moving an application that uses MySQL from on-premises to Google Cloud. The application will run on Compute Engine and will use Cloud SQL. You want to cut over to the Compute Engine deployment of the application with minimal downtime and no data loss to your customers. You want to migrate the application with minimal modification. You also need to determine the cutover strategy. What should you do?

- A. * 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server.* 2. Stop the on-premises application.* 3. Create a mysqldump of the on-premises MySQL server
- B. * 4. Upload the dump to a Cloud Storage bucket.* 5. Import the dump into Cloud SQL.* 6. Modify the source code of the application to write queries to both databases and read from its local database.* 7. Start the Compute Engine application

- C. 8. Stop the on-premises application.
- D. * 1. Set up Cloud SQL proxy and MySQL proxy.* 2. Create a mysqldump of the on-premises MySQL serve
- E. * 3. Upload the dump to a Cloud Storage bucket.* 4. Import the dump into Cloud SQ
- F. * 5. Stop the on-premises applicatio
- G. * 6. Start the Compute Engine application.
- H. * 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL serve
- I. * 2. Stop the on-premises application.* 3. Start the Compute Engine application, configured to read and write to the on-premises MySQL serve
- J. * 4. Create the replication configuration in Cloud SQL.* 5. Configure the source database server to accept connections from the Cloud SQL replic
- K. * 6. Finalize the Cloud SQL replica configuration.* 7. When replication has been completed, stop the Compute Engine applicatio
- L. * 8. Promote the Cloud SQL replica to a standalone instance.* 9. Restart the Compute Engine application, configured to read and write to the Cloud SQL standalone instance.
- M. * 1. Stop the on-premises application.* 2. Create a mysqldump of the on-premises MySQL serve
- N. * 3. Upload the dump to a Cloud Storage bucket.* 4. Import the dump into Cloud SQL.* 5. Start the application on Compute Engine.

Answer: C

Explanation:

External replica promotion migration In the migration strategy of external replica promotion, you create an external database replica and synchronize the existing data to that replica. This can happen with minimal downtime to the existing database. When you have a replica database, the two databases have different roles that are referred to in this document as primary and replica. After the data is synchronized, you promote the replica to be the primary in order to move the management layer with minimal impact to database uptime. In Cloud SQL, an easy way to accomplish the external replica promotion is to use the automated migration workflow. This process automates many of the steps that are needed for this type of migration.

<https://cloud.google.com/architecture/migrating-mysql-to-cloudsql-concept>

- The best option for migrating your MySQL database is to use an external replica promotion. In this strategy, you create a replica database and set your existing database as the primary. You wait until the two databases are in sync, and you then promote your MySQL replica database to be the primary. This process minimizes database downtime related to the database migration. https://cloud.google.com/architecture/migrating-mysql-to-cloudsql-concept#external_replica_promotion_migrat

NEW QUESTION 217

- (Exam Topic 5)

You are configuring the cloud network architecture for a newly created project in Google Cloud that will host applications in Compute Engine. Compute Engine virtual machine instances will be created in two different subnets (sub-a and sub-b) within a single region.

- Instances in sub-a will have public IP addresses
- Instances in sub-b will have only private IP addresses

To download updated packages, instances must connect to a public repository outside the boundaries of Google Cloud. You need to allow sub-b to access the external repository. What should you do?

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

Answer: B

NEW QUESTION 218

- (Exam 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 Logging
- D. Consult logs for Kubernetes Engine and Cloud SQL.
- E. In the GCP Console, navigate to Cloud SQL
- F. Restore the latest backup
- G. Use `kubectl` to restart all pods.

Answer: C

NEW QUESTION 220

- (Exam 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 better 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 that 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 222

- (Exam Topic 5)

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

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

Answer: D

Explanation:

<https://aws.amazon.com/blogs/aws/building-three-tier-architectures-with-security-groups/>

Google Cloud Platform(GCP) enforces firewall rules through rules and tags. GCP rules and tags can be defined once and used across all regions.

References: <https://cloud.google.com/docs/compare/openstack/> <https://aws.amazon.com/it/blogs/aws/building-three-tier-architectures-with-security-groups/>

NEW QUESTION 227

- (Exam Topic 5)

Your company and one of its partners each have a Google Cloud project in separate organizations. Your company's project (prj-a) runs in Virtual Private Cloud (vpc-a). The partner's project (prj-b) runs in vpc-b. There are two instances running on vpc-a and one instance running on vpc-b. Subnets defined in both VPCs are not overlapping. You need to ensure that all instances communicate with each other via internal IPs, minimizing latency and maximizing throughput. What should you do?

- A. Set up a network peering between vpc-a and vpc-b
- B. Set up a VPN between vpc-a and vpc-b using Cloud VPN
- C. Configure IPsec TCP forwarding on the instance in vpc-b and then launch the following gcloud command from one of the instances in vpc-b: gcloud compute start-ipsec-tunnel INSTANCE_NAME --ipsec-key local-host-port:local-host:22
* 1. Create an additional instance in vpc-a * 2. Create an additional instance in vpc-b * 3. Install OpenVPN in newly created instances * 4. Configure a VPN tunnel between vpc-a and vpc-b with the help of OpenVPN

Answer: C

NEW QUESTION 232

- (Exam Topic 5)

One of your primary business objectives is being able to trust the data stored in your application. You want to log all changes to the application data. How can you design your logging system to verify authenticity of your logs?

- A. Write the log concurrently in the cloud and on premises.
- B. Use a SQL database and limit who can modify the log table.
- C. Digitally sign each timestamp and log entry and store the signature.
- D. Create a JSON dump of each log entry and store it in Google Cloud Storage.

Answer: C

Explanation:

<https://cloud.google.com/storage/docs/access-logs>

References: <https://cloud.google.com/logging/docs/reference/tools/gcloud-logging>

NEW QUESTION 233

- (Exam 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 238

- (Exam 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 242

- (Exam 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 for 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 in the different Shared VPCs
- C. Migrate the projects from the acquired company into your company's Google Cloud organization. Relaunch the instances in your company's Shared VPC
- D. Set up a Cloud VPN gateway in each Shared VPC and peer Cloud VPNs

Answer: B

NEW QUESTION 245

- (Exam Topic 5)

Your company provides a recommendation engine for retail customers. You are providing retail customers with an API where they can submit a user ID and the API returns a list of recommendations for that user. You are responsible for the API lifecycle and want to ensure stability for your customers in case the API makes backward-incompatible changes. You want to follow Google-recommended practices. What should you do?

- A. Create a distribution list of all customers to inform them of an upcoming backward-incompatible change at least one month before replacing the old API with the new API.
- B. Create an automated process to generate API documentation, and update the public API documentation as part of the CI/CD process when deploying an update to the API.
- C. Use a versioning strategy for the APIs that increases the version number on every backward-incompatible change.
- D. Use a versioning strategy for the APIs that adds the suffix "DEPRECATED" to the current API version number on every backward-incompatible change.
- E. Use the current version number for the new API.

Answer: C

Explanation:

<https://cloud.google.com/apis/design/versioning>

All Google API interfaces must provide a major version number, which is encoded at the end of the protobuf package, and included as the first part of the URI path for REST APIs. If an API introduces a breaking change, such as removing or renaming a field, it must increment its API version number to ensure that existing user code does not suddenly break.

NEW QUESTION 250

- (Exam 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 251

- (Exam Topic 5)

You need to deploy an application on Google Cloud that must run on a Debian Linux environment. The application requires extensive configuration in order to operate correctly. You want to ensure that you can install Debian distribution updates with minimal manual intervention whenever they become available. What should you do?

- A. Create a Compute Engine instance template using the most recent Debian image
- B. Create an instance from this template, and install and configure the application as part of the startup script
- C. Repeat this process whenever a new Google-managed Debian image becomes available.
- D. Create a Debian-based Compute Engine instance, install and configure the application, and use OS patch management to install available updates.
- E. Create an instance with the latest available Debian image
- F. Connect to the instance via SSH, and install and configure the application on the instance
- G. Repeat this process whenever a new Google-managed Debian image becomes available.

- H. Create a Docker container with Debian as the base image
- I. Install and configure the application as part of the Docker image creation process
- J. Host the container on Google Kubernetes Engine and restart the container whenever a new update is available.

Answer: B

Explanation:

Reference: <https://cloud.google.com/compute/docs/os-patch-management>

NEW QUESTION 253

- (Exam 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 update
- D. Select the Proactive update mode.
- E. Start a new rolling update
- 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 258

- (Exam Topic 5)

You are developing your microservices application on Google Kubernetes Engine. During testing, you want to validate the behavior of your application in case a specific microservice should suddenly crash. What should you do?

- A. Add a taint to one of the nodes of the Kubernetes cluster
- B. For the specific microservice, configure a pod anti-affinity label that has the name of the tainted node as a value.
- C. Use Istio's fault injection on the particular microservice whose faulty behavior you want to simulate.
- D. Destroy one of the nodes of the Kubernetes cluster to observe the behavior.
- E. Configure Istio's traffic management features to steer the traffic away from a crashing microservice.

Answer: B

Explanation:

Microservice runs on all nodes. The Micro service runs on Pod, Pod runs on Nodes. Nodes is nothing but Virtual machines. Once deployed the application microservices will get deployed across all Nodes. Destroying one node may not mimic the behaviour of microservice crashing as it may be running in other nodes.
link: <https://istio.io/latest/docs/tasks/traffic-management/fault-injection/>

NEW QUESTION 260

- (Exam Topic 6)

For this question, refer to the Dress4Win case study. To be legally compliant during an audit, Dress4Win must be able to give insights in all administrative actions that modify the configuration or metadata of resources on Google Cloud.
What should you do?

- A. Use Stackdriver Trace to create a trace list analysis.
- B. Use Stackdriver Monitoring to create a dashboard on the project's activity.
- C. Enable Cloud Identity-Aware Proxy in all projects, and add the group of Administrators as a member.
- D. Use the Activity page in the GCP Console and Stackdriver Logging to provide the required insight.

Answer: A

Explanation:

<https://cloud.google.com/logging/docs/audit/>

NEW QUESTION 262

- (Exam Topic 6)

For this question, refer to the Dress4Win case study. You are responsible for the security of data stored in Cloud Storage for your company, Dress4Win. You have already created a set of Google Groups and assigned the appropriate users to those groups. You should use Google best practices and implement the simplest design to meet the requirements.

Considering Dress4Win's business and technical requirements, what should you do?

- A. Assign custom IAM roles to the Google Groups you created in order to enforce security requirements. Encrypt data with a customer-supplied encryption key when storing files in Cloud Storage.
- B. Assign custom IAM roles to the Google Groups you created in order to enforce security requirements. Enable default storage encryption before storing files in Cloud Storage.
- C. Assign predefined IAM roles to the Google Groups you created in order to enforce security requirements. Utilize Google's default encryption at rest when storing files in Cloud Storage.
- D. Assign predefined IAM roles to the Google Groups you created in order to enforce security requirement
- E. Ensure that the default Cloud KMS key is set before storing files in Cloud Storage.

Answer: D

Explanation:

<https://cloud.google.com/iam/docs/understanding-service-accounts>

NEW QUESTION 263

- (Exam Topic 6)

For this question, refer to the Dress4Win case study. Which of the compute services should be migrated as –is and would still be an optimized architecture for performance in the cloud?

- A. Web applications deployed using App Engine standard environment
- B. RabbitMQ deployed using an unmanaged instance group
- C. Hadoop/Spark deployed using Cloud Dataproc Regional in High Availability mode
- D. Jenkins, monitoring, bastion hosts, security scanners services deployed on custom machine types

Answer: C

NEW QUESTION 265

- (Exam Topic 7)

For this question, refer to the TerramEarth case study.

You start to build a new application that uses a few Cloud Functions for the backend. One use case requires a Cloud Function func_display to invoke another Cloud Function func_query. You want func_query only to accept invocations from func_display. You also want to follow Google's recommended best practices. What should you do?

- A. Create a token and pass it in as an environment variable to func_display
- B. When invoking func_query, include the token in the request Pass the same token to func_query and reject the invocation if the tokens are different.
- C. Make func_query 'Require authentication.' Create a unique service account and associate it to func_display
- D. Grant the service account invoker role for func_query
- E. Create an id token in func_display and include the token to the request when invoking func_query.
- F. Make func_query 'Require authentication' and only accept internal traffic
- G. Create those two functions in the same VP
- H. Create an ingress firewall rule for func_query to only allow traffic from func_display.
- I. Create those two functions in the same project and VP
- J. Make func_query only accept internal traffic. Create an ingress firewall for func_query to only allow traffic from func_display
- K. Also, make sure both functions use the same service account.

Answer: B

Explanation:

https://cloud.google.com/functions/docs/securing/authenticating#authenticating_function_to_function_calls

NEW QUESTION 269

- (Exam Topic 7)

TerramEarth has a legacy web application that you cannot migrate to cloud. However, you still want to build a cloud-native way to monitor the application. If the application goes down, you want the URL to point to a "Site is unavailable" page as soon as possible. You also want your Ops team to receive a notification for the issue. You need to build a reliable solution for minimum cost

What should you do?

- A. Create a scheduled job in Cloud Run to invoke a container every minute
- B. The container will check the application URL If the application is down, switch the URL to the "Site is unavailable" page, and notify the Ops team.
- C. Create a cron job on a Compute Engine VM that runs every minute
- D. The cron job invokes a Python program to check the application URL If the application is down, switch the URL to the "Site is unavailable" page, and notify the Ops team.
- E. Create a Cloud Monitoring uptime check to validate the application URL If it fails, put a message in a Pub/Sub queue that triggers a Cloud Function to switch the URL to the "Site is unavailable" page, and notify the Ops team.
- F. Use Cloud Error Reporting to check the application URL If the application is down, switch the URL to the "Site is unavailable" page, and notify the Ops team.

Answer: C

Explanation:

<https://cloud.google.com/blog/products/management-tools/how-to-use-pubsub-as-a-cloud-monitoring-notification>

NEW QUESTION 272

- (Exam Topic 7)

For this question, refer to the TerramEarth case study. A new architecture that writes all incoming data to BigQuery has been introduced. You notice that the data is dirty, and want to ensure data quality on an automated daily basis while managing cost.

What should you do?

- A. Set up a streaming Cloud Dataflow job, receiving data by the ingestion process
- B. Clean the data in a Cloud Dataflow pipeline.
- C. Create a Cloud Function that reads data from BigQuery and cleans it
- D. Trigger it
- E. Trigger the Cloud Function from a Compute Engine instance.
- F. Create a SQL statement on the data in BigQuery, and save it as a view
- G. Run the view daily, and save the result to a new table.
- H. Use Cloud Dataprep and configure the BigQuery tables as the source
- I. Schedule a daily job to clean the data.

Answer: A

NEW QUESTION 277

- (Exam Topic 7)

You have broken down a legacy monolithic application into a few containerized RESTful microservices. You want to run those microservices on Cloud Run. You also want to make sure the services are highly available with low latency to your customers. What should you do?

- A. Deploy Cloud Run services to multiple availability zone
- B. Create Cloud Endpoints that point to the service
- C. Create a global HTTP(S) Load Balancing instance and attach the Cloud Endpoints to its backend.
- D. Deploy Cloud Run services to multiple regions Create serverless network endpoint groups pointing to the service
- E. Add the serverless NEGs to a backend service that is used by a global HTTP(S) Load Balancing instance.
- F. Cloud Run services to multiple region
- G. In Cloud DNS, create a latency-based DNS name that points to the services.
- H. Deploy Cloud Run services to multiple availability zone
- I. Create a TCP/IP global load balance
- J. Add the Cloud Run Endpoints to its backend service.

Answer: B

Explanation:

<https://cloud.google.com/run/docs/multiple-regions>

NEW QUESTION 282

- (Exam Topic 7)

You are migrating a Linux-based application from your private data center to Google Cloud. The TerraEarth security team sent you several recent Linux vulnerabilities published by Common Vulnerabilities and Exposures (CVE). You need assistance in understanding how these vulnerabilities could impact your migration. What should you do?

- A. Open a support case regarding the CVE and chat with the support engineer.
- B. Read the CVEs from the Google Cloud Status Dashboard to understand the impact.
- C. Read the CVEs from the Google Cloud Platform Security Bulletins to understand the impact
- D. Post a question regarding the CVE in Stack Overflow to get an explanation
- E. Post a question regarding the CVE in a Google Cloud discussion group to get an explanation

Answer: AC

Explanation:

<https://cloud.google.com/support/bulletins>

NEW QUESTION 286

- (Exam Topic 7)

For this question, refer to the TerraEarth case study. To be compliant with European GDPR regulation, TerraEarth is required to delete data generated from its European customers after a period of 36 months when it contains personal data. In the new architecture, this data will be stored in both Cloud Storage and BigQuery. What should you do?

- A. Create a BigQuery table for the European data, and set the table retention period to 36 month
- B. For Cloud Storage, use gsutil to enable lifecycle management using a DELETE action with an Age condition of 36 months.
- C. Create a BigQuery table for the European data, and set the table retention period to 36 month
- D. For Cloud Storage, use gsutil to create a SetStorageClass to NONE action when with an Age condition of 36 months.
- E. Create a BigQuery time-partitioned table for the European data, and set the partition expiration period to 36 month
- F. For Cloud Storage, use gsutil to enable lifecycle management using a DELETE action with an Age condition of 36 months.
- G. Create a BigQuery time-partitioned table for the European data, and set the partition period to 36 month
- H. For Cloud Storage, use gsutil to create a SetStorageClass to NONE action with an Age condition of 36 months.

Answer: C

Explanation:

<https://cloud.google.com/bigquery/docs/managing-partitioned-tables#partition-expiration> <https://cloud.google.com/storage/docs/lifecycle>

NEW QUESTION 287

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

Professional-Cloud-Architect Practice Exam Features:

- * Professional-Cloud-Architect Questions and Answers Updated Frequently
- * Professional-Cloud-Architect Practice Questions Verified by Expert Senior Certified Staff
- * Professional-Cloud-Architect Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * Professional-Cloud-Architect Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The Professional-Cloud-Architect Practice Test Here](#)