


Professional-Cloud-Architect Dumps

Google Certified Professional - Cloud Architect (GCP)

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



NEW QUESTION 1

- (Topic 1)

For this question, refer to the Mountkirk Games case study

Mountkirk Games needs to create a repeatable and configurable mechanism for deploying isolated application environments. Developers and testers can access each other's environments and resources, but they cannot access staging or production resources. The staging environment needs access to some services from production.

What should you do to isolate development environments from staging and production?

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

Answer: D

NEW QUESTION 2

- (Topic 1)

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

Mountkirk Games wants to set up a real-time analytics platform for their new game. The new platform must meet their technical requirements. Which combination of Google technologies will meet all of their requirements?

- A. Container Engine, Cloud Pub/Sub, and Cloud SQL
- B. Cloud Dataflow, Cloud Storage, Cloud Pub/Sub, and BigQuery
- C. Cloud SQL, Cloud Storage, Cloud Pub/Sub, and Cloud Dataflow
- D. Cloud Dataproc, Cloud Pub/Sub, Cloud SQL, and Cloud Dataflow
- E. Cloud Pub/Sub, Compute Engine, Cloud Storage, and Cloud Dataproc

Answer: B

Explanation:

A real time requires Stream / Messaging so Pub/Sub, Analytics by Big Query.

Ingest millions of streaming events per second from anywhere in the world with Cloud Pub/Sub, powered by Google's unique, high-speed private network. Process the streams with Cloud Dataflow to ensure reliable, exactly-once, low-latency data transformation. Stream the transformed data into BigQuery, the cloud-native data warehousing service, for immediate analysis via SQL or popular visualization tools.

From scenario: They plan to deploy the game's backend on Google Compute Engine so they can capture streaming metrics, run intensive analytics.

Requirements for Game Analytics Platform

- ? Dynamically scale up or down based on game activity
- ? Process incoming data on the fly directly from the game servers
- ? Process data that arrives late because of slow mobile networks
- ? Allow SQL queries to access at least 10 TB of historical data
- ? Process files that are regularly uploaded by users' mobile devices
- ? Use only fully managed services

References: <https://cloud.google.com/solutions/big-data/stream-analytics/>

NEW QUESTION 3

- (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 4

- (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 5

- (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.-1712523161.1522785064

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

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

From scenario: Business Requirements

Decrease unplanned vehicle downtime to less than 1 week, without increasing the cost of carrying surplus inventory

Support the dealer network with more data on how their customers use their equipment to better position new products and services

Have the ability to partner with different companies – especially with seed and fertilizer suppliers in the fast-growing agricultural business – to create compelling joint offerings for their customers.

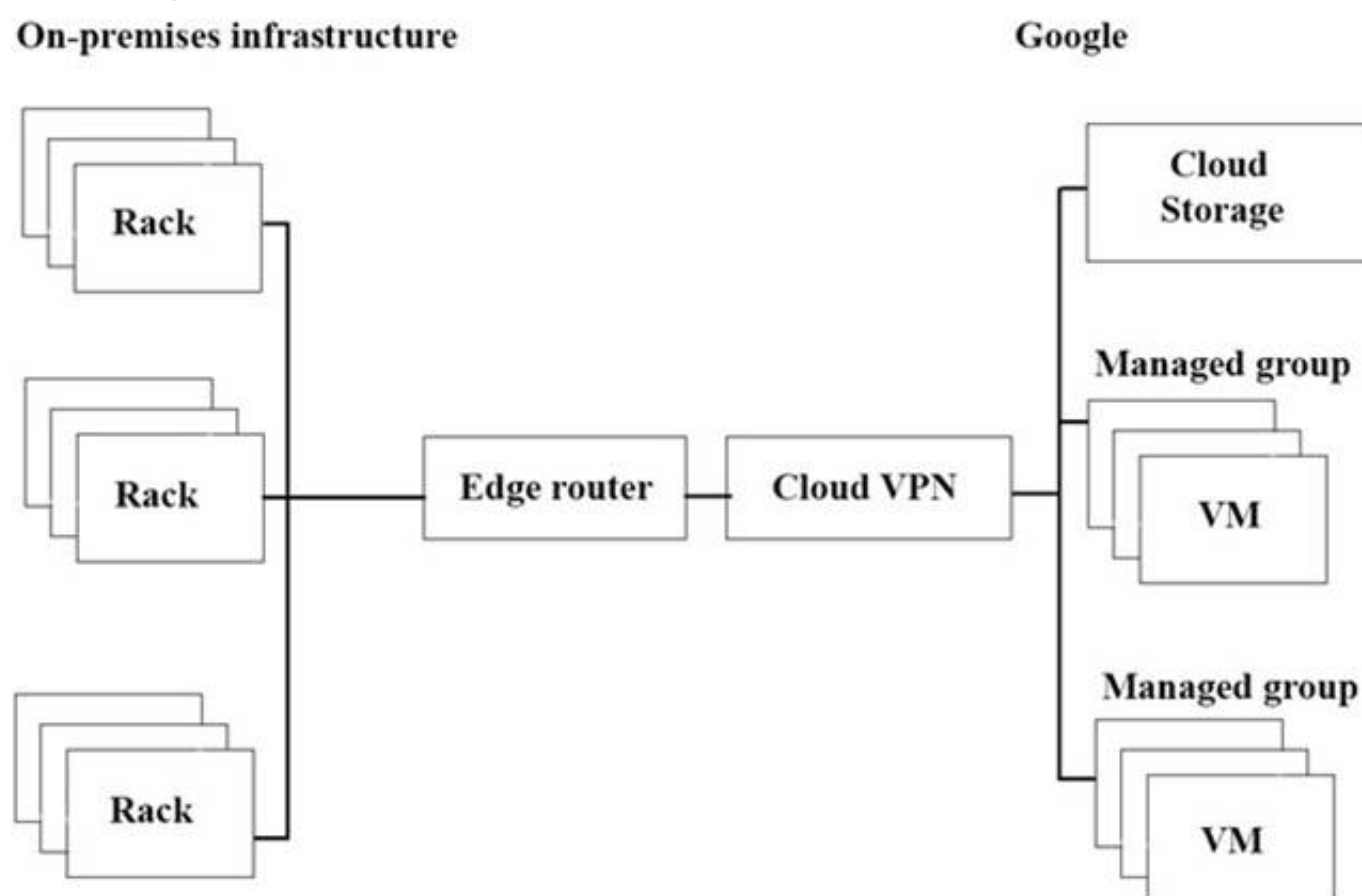
Reference: <https://cloud.google.com/certification/guides/cloud-architect/casestudy-terramearth>

NEW QUESTION 6

- (Topic 3)

For this question, refer to the JencoMart case study.

The migration of JencoMart's application to Google Cloud Platform (GCP) is progressing too slowly. The infrastructure is shown in the diagram. You want to maximize throughput. What are three potential bottlenecks? (Choose 3 answers.)



- A. A single VPN tunnel, which limits throughput
- B. A tier of Google Cloud Storage that is not suited for this task
- C. A copy command that is not suited to operate over long distances
- D. Fewer virtual machines (VMs) in GCP than on-premises machines
- E. A separate storage layer outside the VMs, which is not suited for this task
- F. Complicated internet connectivity between the on-premises infrastructure and GCP

Answer: ADF

NEW QUESTION 7

- (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 8

- (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 9

- (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
- E. 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 10

- (Topic 4)

For this question, refer to the Dress4Win case study.

As part of Dress4Win's plans to migrate to the cloud, they want to be able to set up a managed logging and monitoring system so they can handle spikes in their traffic load.

They want to ensure that:

- The infrastructure can be notified when it needs to scale up and down to handle the ebb and flow of usage throughout the day
- Their administrators are notified automatically when their application reports errors.
- They can filter their aggregated logs down in order to debug one piece of the application across many hosts

Which Google StackDriver features should they use?

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

Answer: D

NEW QUESTION 10

- (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 14

- (Topic 5)

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

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

Answer: A

NEW QUESTION 17

- (Topic 5)

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

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

Answer: A

Explanation:

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

? <https://cloud.google.com/architecture/framework/system-design>

? <https://cloud.google.com/architecture/identity/best-practices-for-planning>

? <https://cloud.google.com/resource-manager/docs/creating-managing-folders>

NEW QUESTION 22

- (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

A.

Answer: B

NEW QUESTION 26

- (Topic 5)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 29

- (Topic 5)

Your company is planning to upload several important files to Cloud Storage. After the upload is completed, they want to verify that the upload content is identical to what they have on- premises. You want to minimize the cost and effort of performing this check. What should you do?

- A.
- 1) Use gsutil -m to upload all the files to Cloud Storage.
 - 2) Use gsutil cp to download the uploaded files
 - 3) Use Linux diff to compare the content of the files
- B.
- 1) Use gsutil -m to upload all the files to Cloud Storage.
 - 2) Develop a custom Java application that computes CRC32C hashes
 - 3) Use gsutil ls -L gs://[YOUR_BUCKET_NAME] to collect CRC32C hashes of the uploaded files
 - 4) Compare the hashes
- C.
- 1) Use Linux shasum to compute a digest of files you want to upload
 - 2) Use gsutil -m to upload all the files to the Cloud Storage
 - 3) Use gsutil cp to download the uploaded files
 - 4) Use Linux shasum to compute a digest of the downloaded files 5. Compare the hashes
- D.
- 1) Use gsutil -m to upload all the files to Cloud Storage.
 - 2) Use gsutil hash -c FILE_NAME to generate CRC32C hashes of all on-premises files 3) Use gsutil ls -L gs://[YOUR_BUCKET_NAME] to collect CRC32C hashes of the uploaded files
 - 4) Compare the hashes

A.

Answer: D

Explanation:

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

NEW QUESTION 31

- (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: D

Explanation:

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

NEW QUESTION 34

- (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 protects 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 36

- (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 hours
- F. Authenticate users via Cloud Identity.
- G. Create an App Engine web application where users can upload images for the next 24 hours
- H. Authenticate users via Cloud Identity.

Answer: A

Explanation:

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

NEW QUESTION 37

- (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 want 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 40

- (Topic 5)

A development manager is building a new application. He asks you to review his requirements and identify what cloud technologies he can use to meet them. The application must

- * 1. Be based on open-source technology for cloud portability
- * 2. Dynamically scale compute capacity based on demand
- * 3. Support continuous software delivery
- * 4. Run multiple segregated copies of the same application stack
- * 5. Deploy application bundles using dynamic templates
- * 6. Route network traffic to specific services based on URL

Which combination of technologies will meet all of his requirements?

- A. Google Container Engine, Jenkins, and Helm
- B. Google Container Engine and Cloud Load Balancing
- C. Google Compute Engine and Cloud Deployment Manager
- D. Google Compute Engine, Jenkins, and Cloud Load Balancing

Answer: A

Explanation:

Helm for managing Kubernetes

Kubernetes can base on the URL to route traffic to different location (path)

<https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer> eg. apiVersion: networking.k8s.io/v1beta1

kind: Ingress metadata:

name: fanout-ingress spec:

rules:

- http: paths:

- path: /* backend:

serviceName: web servicePort: 8080

- path: /v2/* backend: serviceName: web2 servicePort: 8080

NEW QUESTION 43

- (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 48

- (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 52

- (Topic 5)

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

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

Answer: C

Explanation:

From: <https://cloud.google.com/architecture/best-practices-for-building-containers>

Under: Tagging using the Git commit hash (bottom of page almost)

"In this case, a common way of handling version numbers is to use the Git commit SHA-1 hash (or a short version of it) as the version number. By design, the Git commit hash is immutable and references a specific version of your software.

You can use this commit hash as a version number for your software, but also as a tag for the Docker image built from this specific version of your software. Doing so makes Docker images traceable: because in this case the image tag is immutable, you instantly know which specific version of your software is running inside a given container."

NEW QUESTION 57

- (Topic 5)

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

- A. Configure Private Google Access for on-premises only.
- B. Perform the following tasks: 1) Create a service account. 2) Give the BigQuery JobUser role and Storage Reader role to the service account. 3) Remove all other IAM access from the project.
- C. Configure VPC Service Controls and configure Private Google Access.
- D. Configure Private Google Access.

Answer: C

Explanation:

<https://cloud.google.com/vpc-service-controls/docs/overview>

VPC Service Controls improves your ability to mitigate the risk of data exfiltration from Google Cloud services such as Cloud Storage and BigQuery.

NEW QUESTION 60

- (Topic 5)

Your BigQuery project has several users. For audit purposes, you need to see how many queries each user ran in the last month.

- A. Connect Google Data Studio to BigQuer
- B. Create a dimension for the users and a metric for the amount of queries per user.
- C. In the BigQuery interface, execute a query on the JOBS table to get the required information.
- D. Use 'bq show' to list all job
- E. Per job, use 'bq ls' to list job information and get the required information.
- F. Use Cloud Audit Logging to view Cloud Audit Logs, and create a filter on the query operation to get therequired information.

Answer: C

Explanation:

<https://cloud.google.com/bigquery/docs/managing-jobs>

NEW QUESTION 65

- (Topic 5)

Your company has a Google Cloud project that uses BigQuery for data warehousing on a pay-per-use basis. You want to monitor queries in real time to discover the most costly queries and which users spend the most. What should you do?

- A.
 - * 1. Create a Cloud Logging sink to export BigQuery data access logs to Cloud Storage.
 - * 2. Develop a Dataflow pipeline to compute the cost of queries split by users.
- B.
 - * 1. Create a Cloud Logging sink to export BigQuery data access logs to BigQuery.
 - * 2. Perform a BigQuery query on the generated table to extract the information you need.
- C.
 - * 1. Activate billing export into BigQuery.
 - * 2. Perform a BigQuery query on the billing table to extract the information you need.
- D.
 - * 1. In the BigQuery dataset that contains all the tables to be queried, add a label for each user that can launch a query.
 - * 2. Open the Billing page of the project.
 - * 3. Select Reports.
 - * 4. Select BigQuery as the product and filter by the user you want to check.

A.

Answer: C

Explanation:

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

NEW QUESTION 69

- (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 70

- (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 cluste
- 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 75

- (Topic 5)

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

- A. Set the memcache service level to dedicate
- B. Create a key from the hash of the query, and return database values from memcache before issuing a query to Cloud SQL.
- C. Set the memcache service level to dedicate
- D. Create a cron task that runs every minute to populate the cache with keys containing query results.
- E. Set the memcache service level to share
- F. Create a cron task that runs every minute to save all expected queries to a key called “cached-queries”.
- G. Set the memcache service level to share
- H. Create a key called “cached-queries”, and return 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 78

- (Topic 5)

You are creating a solution to remove backup files older than 90 days from your backup Cloud Storage bucket. You want to optimize ongoing Cloud Storage spend. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 80

- (Topic 5)

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

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

Answer: C

Explanation:

It is time-series data, So Big Table. <https://cloud.google.com/bigtable/docs/schema-design-time-series>

Google Cloud Bigtable is a scalable, fully-managed NoSQL wide-column database that is suitable for both real-time access and analytics workloads.

Good for:

- ? Low-latency read/write access
- ? High-throughput analytics
- ? Native time series support
- ? Common workloads:
- ? IoT, finance, adtech
- ? Personalization, recommendations
- ? Monitoring
- ? Geospatial datasets
- ? Graphs

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

NEW QUESTION 82

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 86

- (Topic 5)

Your customer runs a web service used by e-commerce sites to offer product recommendations to users. The company has begun experimenting with a machine learning model on Google Cloud Platform to improve the quality of results. What should the customer do to improve their model's results over time?

- A. Export Cloud Machine Learning Engine performance metrics from Stackdriver to BigQuery, to be used to analyze the efficiency of the model.
- B. Build a roadmap to move the machine learning model training from Cloud GPUs to Cloud TPUs, which offer better results.
- C. Monitor Compute Engine announcements for availability of newer CPU architectures, and deploy the model to them as soon as they are available for additional performance.
- D. Save a history of recommendations and results of the recommendations in BigQuery, to be used as training data.

Answer: D

Explanation:

<https://cloud.google.com/solutions/building-a-serverless-ml-model>

NEW QUESTION 89

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 90

- (Topic 5)

You deploy your custom Java application to Google App Engine. It fails to deploy and gives you the following stack trace.

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

What should you do?

- A. Upload missing JAR files and redeploy your application.
- B. Digitally sign all of your JAR files and redeploy your application

C. Recompile the CLoakedServlet class using and MD5 hash instead of SHA1

Answer: B

NEW QUESTION 94

- (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 95

- (Topic 5)

Your company has an application that is running on multiple instances of Compute Engine. It generates 1 TB per day of logs. For compliance reasons, the logs need to be kept for at least two years. The logs need to be available for active query for 30 days. After that, they just need to be retained for audit purposes. You want to implement a storage solution that is compliant, minimizes costs, and follows Google-recommended practices. What should you do?

- A.
 - * 1. Install the Cloud Ops agent on all instances.
 - * 2. Create a sink to export logs into a partitioned BigQuery table.
 - * 3. Set a time_partitioning_expiration of 30 days.
- B.
 - * 1. Install the Cloud Ops agent on all instances.
 - * 2. Create a sink to export logs into a regional Cloud Storage bucket.
 - * 3. Create an Object Lifecycle rule to move files into a Coldline Cloud Storage bucket after one month.
 - * 4. Configure a retention policy at the bucket level to create a lock.
- C.
 - * 1. Create a daily cron job, running on all instances, that uploads logs into a partitioned BigQuery table.
 - * 2. Set a time_partitioning_expiration of 30 days.
- D.
 - * 1. Write a daily cron job, running on all instances, that uploads logs into a Cloud Storage bucket.
 - * 2. Create a sink to export logs into a regional Cloud Storage bucket.
 - * 3. Create an Object Lifecycle rule to move files into a Coldline Cloud Storage bucket after one month.

A.

Answer: B

Explanation:

The practice for managing logs generated on Compute Engine on Google Cloud is to install the Cloud Logging agent and send them to Cloud Logging. The sent logs will be aggregated into a Cloud Logging sink and exported to Cloud Storage. The reason for using Cloud Storage as the destination for the logs is that the requirement in question requires setting up a lifecycle based on the storage period. In this case, the log will be used for active queries for 30 days after it is saved, but after that, it needs to be stored for a longer period of time for auditing purposes. If the data is to be used for active queries, we can use BigQuery's Cloud Storage data query feature and move the data past 30 days to Coldline to build a cost-optimal solution.

Therefore, the correct answer is as follows

- * 1. Install the Cloud Logging agent on all instances.
- Create a sync that exports the logs to the region's Cloud Storage bucket.
- * 3. Create an Object Lifecycle rule to move the files to the Coldline Cloud Storage bucket after one month.
 - * 4. set up a bucket-level retention policy using bucket locking."

NEW QUESTION 96

- (Topic 5)

Your organization requires that metrics from all applications be retained for 5 years for future analysis in possible legal proceedings. Which approach should you use?

- A. Grant the security team access to the logs in each Project.
- B. Configure Stackdriver Monitoring for all Projects, and export to BigQuery.
- C. Configure Stackdriver Monitoring for all Projects with the default retention policies.
- D. Configure Stackdriver Monitoring for all Projects, and export to Google Cloud Storage.

Answer: D

Explanation:

Overview of storage classes, price, and use cases <https://cloud.google.com/storage/docs/storage-classes>

Why export logs? <https://cloud.google.com/logging/docs/export/>

StackDriver Quotas and Limits for Monitoring <https://cloud.google.com/monitoring/quotas> The BigQuery pricing. <https://cloud.google.com/bigquery/pricing>

NEW QUESTION 99

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 104

- (Topic 5)

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

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

Answer: AD

Explanation:

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

NEW QUESTION 109

- (Topic 5)

You need to design a solution for global load balancing based on the URL path being requested. You need to ensure operations reliability and end-to-end in-transit encryption based on Google best practices. What should you do?

- A. Create a cross-region load balancer with URL Maps.
- B. Create an HTTPS load balancer with URL maps.
- C. Create appropriate instance groups and instance
- D. Configure SSL proxy load balancing.
- E. Create a global forwarding rule.
- F. Configure SSL proxy balancing.

Answer: B

Explanation:

Reference <https://cloud.google.com/load-balancing/docs/https/url-map>

NEW QUESTION 113

- (Topic 5)

Your marketing department wants to send out a promotional email campaign. The development team wants to minimize direct operation management. They project a wide range of possible customer responses, from 100 to 500,000 click-throughs per day. The link leads to a simple website that explains the promotion and collects user information and preferences. Which infrastructure should you recommend? (CHOOSE TWO)

- A. Use Google App Engine to serve the website and Google Cloud Datastore to store user data.
- B. Use a Google Container Engine cluster to serve the website and store data to persistent disk.
- C. Use a managed instance group to serve the website and Google Cloud Bigtable to store user data.
- D. Use a single compute Engine virtual machine (VM) to host a web server, backed by Google Cloud SQL.

Answer: AC

Explanation:

Reference: <https://cloud.google.com/storage-options/>

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

NEW QUESTION 115

- (Topic 5)

You need to deploy a stateful workload on Google Cloud. The workload can scale horizontally, but each instance needs to read and write to the same POSIX filesystem. At high load, the stateful workload needs to support up to 100 MB/s of writes. What should you do?

- A. Use a persistent disk for each instance.
- B. Use a regional persistent disk for each instance.
- C. Create a Cloud Filestore instance and mount it in each instance.
- D. Create a Cloud Storage bucket and mount it in each instance using gcsfuse.

Answer: C

Explanation:

<https://cloud.google.com/storage/docs/gcs-fuse#notes>

Cloud Filestore: Cloud Filestore is a scalable and highly available shared file service fully managed by Google. Cloud Filestore provides persistent storage ideal for shared workloads. It is best suited for enterprise applications requiring persistent, durable, shared storage which is accessed by NFS or requires a POSIX compliant file system.

Reference: <https://cloud.google.com/storage/docs/gcs-fuse>

NEW QUESTION 116

- (Topic 5)

Your company creates rendering software which users can download from the company website. Your company has customers all over the world. You want to minimize latency for all your customers. You want to follow Google-recommended practices. How should you store the files?

- A. Save the files in a Multi-Regional Cloud Storage bucket.
- B. Save the files in a Regional Cloud Storage bucket, one bucket per zone of the region.
- C. Save the files in multiple Regional Cloud Storage buckets, one bucket per zone per region.
- D. Save the files in multiple Multi-Regional Cloud Storage buckets, one bucket per multi- region.

Answer: A

Explanation:

<https://cloud.google.com/storage/docs/locations#location-mr>

NEW QUESTION 121

- (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 122

- (Topic 5)

Your company sends all Google Cloud logs to Cloud Logging. Your security team wants to monitor the logs. You want to ensure that the security team can react quickly if an anomaly such as an unwanted firewall change or server breach is detected. You want to follow Google-recommended practices. What should you do?

- A. Schedule a cron job with Cloud Schedule
- B. The scheduled job queries the logs every minute for the relevant events.
- C. Export logs to BigQuery, and trigger a query in BigQuery to process the log data for the relevant events.
- D. Export logs to a Pub/Sub topic, and trigger Cloud Function with the relevant log events.
- E. Export logs to a Cloud Storage bucket, and trigger Cloud Run with the relevant log events.

Answer: C

Explanation:

<https://cloud.google.com/blog/products/management-tools/automate-your-response-to-a-cloud-logging-event>

NEW QUESTION 124

- (Topic 5)

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

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

Answer: D

Explanation:

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

NEW QUESTION 127

- (Topic 5)

You deploy your custom java application to google app engine. It fails to deploy and gives you the following stack trace:

```
Java.lang.SecurityException : SHA1 digest

At com.google.appengine.runtime.Request.pro

At

Sun.securityutil.manifestEntryVerifier.ver

At java . net . URLClassLoader . defineCla

At sun . reflect . GeneratedMethodAccessors

At

Sun.reflect . DelegatingMethodAccesorImpl.

At java . lang . reflect . MThod . invoke
```

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

Answer: B

NEW QUESTION 130

- (Topic 5)

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

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

Answer: ACE

Explanation:

<https://www.flexera.com/blog/cloud/2013/12/google-compute-engine-live-migration-passes-the-test/>

"With live migration, the virtual machines are moved without any downtime or noticeable service degradation"

NEW QUESTION 134

- (Topic 5)

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

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

Answer: C

NEW QUESTION 137

- (Topic 5)

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

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

Answer: A

NEW QUESTION 140

- (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 IAP 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-tcp-tunnel INSTANCE_NAME_VPC_B 22 --local-host-port=localhost: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 143

- (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 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 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 148

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 152

- (Topic 5)

You have an outage in your Compute Engine managed instance group: all instances 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 156

- (Topic 5)

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

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

Answer: B

Explanation:

Relational databases were not designed to cope with the scale and agility challenges that face modern applications, nor were they built to take advantage of the commodity storage and processing power available today.

NoSQL fits well for:

_ Developers are working with applications that create massive volumes of new, rapidly changing data types — structured, semi-structured, unstructured and polymorphic data.

NEW QUESTION 157

- (Topic 5)

You have created several preemptible Linux virtual machine instances using Google Compute Engine. You want to properly shut down your application before the virtual machines are preempted. What should you do?

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

Answer: C

NEW QUESTION 161

- (Topic 5)

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

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

Answer: D

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/multi-cluster-ingress#how_works

<https://github.com/GoogleCloudPlatform/k8s-multicloud-ingress> <https://cloud.google.com/blog/products/gcp/how-to-deploy-geographically-distributed-services-on-kubernetes-engine-with-kubemci>

NEW QUESTION 165

- (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 166

- (Topic 5)

Your company acquired a healthcare startup and must retain its customers' medical information for up to 4 more years, depending on when it was created. Your corporate policy is to securely retain this data, and then delete it as soon as regulations allow.

Which approach should you take?

- A. Store the data in Google Drive and manually delete records as they expire.
- B. Anonymize the data using the Cloud Data Loss Prevention API and store it indefinitely.
- C. Store the data using the Cloud Storage and use lifecycle management to delete files when they expire.
- D. Store the data in Cloud Storage and run a nightly batch script that deletes all expired data.

Answer: C

Explanation:

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

NEW QUESTION 170

- (Topic 5)

A small number of API requests to your microservices-based application take a very long time. You know that each request to the API can traverse many services. You want to know which service takes the longest in those cases. What should you do?

- A. Set timeouts on your application so that you can fail requests faster.
- B. Send custom metrics for each of your requests to Stackdriver Monitoring.
- C. Use Stackdriver Monitoring to look for insights that show when your API latencies are high.
- D. Instrument your application with Stackdriver Trace in order to break down the request latencies at each microservice.

Answer: D

Explanation:

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

NEW QUESTION 174

- (Topic 5)

You have been asked to select the storage system for the click-data of your company's large portfolio of websites. This data is streamed in from a custom website analytics package at a typical rate of 6,000 clicks per minute, with bursts of up to 8,500 clicks per second. It must be stored for future analysis by your data science and user experience teams. Which storage infrastructure should you choose?

- A. Google Cloud SQL
- B. Google Cloud Bigtable
- C. Google Cloud Storage
- D. Google cloud Datastore

Answer: C

Explanation:

<https://cloud.google.com/bigquery/docs/loading-data-cloud-storage>

NEW QUESTION 178

- (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 183

- (Topic 5)

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

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

Answer: A

Explanation:

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

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

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

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

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

References: <https://cloud.google.com/bigquery/>

NEW QUESTION 185

- (Topic 5)

You want to automate the creation of a managed instance group and a startup script to install the OS package dependencies. You want to minimize the startup time for VMs in the instance group.

What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 186

- (Topic 5)

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

Which of the following best reflects your recommendations for a cost-effective storage allocation?

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

Answer: D

Explanation:

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

NEW QUESTION 188

- (Topic 5)

A news teed web service has the following code running on Google App Engine. During peak load, users report that they can see news articles they already viewed. What is the most likely cause of this problem?

```
import news
from flask import Flask, redirect, request
from flask.ext.api import status
from google.appengine.api import users

app = Flask(__name__)
sessions = {}

@app.route("/")
def homepage():
    user = users.get_current_user()
    if not user:
        return "Invalid login",
        status.HTTP_401_UNAUTHORIZED

    if user not in sessions:
        sessions[user] = {"viewed": []}

    news_articles = news.get_new_news (user, sessions [user]
["viewed"])
    sessions [user] ["viewed"] += [n["id"] for n
in news_articles]

    return news.render(news_articles)

if __name__ == "__main__":
    app.run()
```

- A. The session variable is local to just a single instance.
- B. The session variable is being overwritten in Cloud Datastore.
- C. The URL of the API needs to be modified to prevent caching.
- D. The HTTP Expires header needs to be set to -1 to stop caching.

Answer: A

Explanation:

<https://stackoverflow.com/questions/3164280/google-app-engine-cache-list-in-session-variable?rq=1>

NEW QUESTION 190

- (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 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 193

- (Topic 5)

You have an App Engine application that needs to be updated. You want to test the update with production traffic before replacing the current application version. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 198

- (Topic 5)

You need to develop procedures to verify resilience of disaster recovery for remote recovery using GCP. Your production environment is hosted on-premises. You

need to establish a secure, redundant connection between your on premises network and the GCP network.
What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 202

- (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 207

- (Topic 5)

Your company has an enterprise application running on Compute Engine that requires high availability and high performance. The application has been deployed on two instances in two zones in the same region in active passive mode. The application writes data to a persistent disk in the case of a single zone outage that data should be immediately made available to the other instance in the other zone. You want to maximize performance while minimizing downtime and data loss. What should you do?

- A.
 - * 1. Attach a persistent SSD disk to the first instance
 - * 2. Create a snapshot every hour
 - * 3. In case of a zone outage, recreate a persistent SSD disk in the second instance where data is coming from the created snapshot
- B.
 - * 1 Create a Cloud Storage bucket
 - * 2. Mount the bucket into the first instance with gcs-fuse
 - * 3. In case of a zone outage, mount the Cloud Storage bucket to the second instance with gcs-fuse
- C.
 - * 1 Attach a local SSD to the first instance disk
 - * 2. Execute an rsync command every hour where the target is a persistent SSD disk attached to the second instance
 - * 3. In case of a zone outage, use the second instance
- D.
 - * 1. Attach a regional SSD persistent disk to the first instance
 - * 2. In case of a zone outage, force-attach the disk to the other instance

A.

Answer: D

NEW QUESTION 208

- (Topic 5)

A lead engineer wrote a custom tool that deploys virtual machines in the legacy data center. He wants to migrate the custom tool to the new cloud environment. You want to advocate for the adoption of Google Cloud Deployment Manager. What are two business risks of migrating to Cloud Deployment Manager? Choose 2 answers.

- A. Cloud Deployment Manager uses Python.
- B. Cloud Deployment Manager APIs could be deprecated in the future.
- C. Cloud Deployment Manager is unfamiliar to the company's engineers.
- D. Cloud Deployment Manager requires a Google APIs service account to run.
- E. Cloud Deployment Manager can be used to permanently delete cloud resources.
- F. Cloud Deployment Manager only supports automation of Google Cloud resources.

Answer: CF

Explanation:

<https://cloud.google.com/deployment-manager/docs/deployments/deleting-deployments>

NEW QUESTION 213

- (Topic 5)

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

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

Answer: B

NEW QUESTION 218

- (Topic 5)

You are managing several internal applications that are deployed on Compute Engine. Business users inform you that an application has become very slow over the past few days. You want to find the underlying cause in order to solve the problem. What should you do first?

- A. Inspect the logs and metrics from the instances in Cloud Logging and Cloud Monitoring.
- B. Restore a backup of the application database from a time before the application became slow.
- C. Deploy the applications on a managed instance group with autoscaling enable
- D. Add a load balancer in front of the managed instance group, and have the users connect to the IP of the load balancer.
- E. Change the Compute Engine Instances behind the application to a machine type with more CPU and memory.

Answer: A

Explanation:

When an application becomes slow, the first step you should take is to gather information about the underlying cause of the problem. One way to do this is by inspecting the logs and metrics from the instances where the application is deployed. Google Cloud Platform (GCP) provides tools such as Cloud Logging and Cloud Monitoring that can help you to collect and analyze this information. By reviewing the logs and metrics from the instances, you may be able to identify issues such as resource shortages (e.g. CPU, memory, or disk), network problems, or application errors that are causing the performance issues. Once you have identified the underlying cause of the problem, you can take steps to resolve it.

NEW QUESTION 219

- (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 published 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 than 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 220

- (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 224

- (Topic 5)

Your company has an application running on a deployment in a GKE cluster. You have a separate cluster for development, staging and production. You have discovered that the team is able to deploy a Docker image to the production cluster without first testing the deployment in development and then staging. You want to allow the team to have autonomy but want to prevent this from happening. You want a Google Cloud solution that can be implemented quickly with minimal effort. What should you do?

- A. Create a Kubernetes admission controller to prevent the container from starting if it is not approved for usage in the given environment
- B. Configure a Kubernetes lifecycle hook to prevent the container from starting if it is not approved for usage in the given environment
- C. Implement a corporate policy to prevent teams from deploying Docker image to an environment unless the Docker image was tested in an earlier environment
- D. Configure the binary authorization policies for the development, staging and production cluster
- E. Create attestations as part of the continuous integration pipeline

Answer: D

Explanation:

<https://cloud.google.com/architecture/prep-kubernetes-engine-for-prod#binary-authorization>

The most common Binary Authorization use cases involve attestations. An attestation certifies that a specific image has completed a previous stage, as described previously. You configure the Binary Authorization policy to verify the attestation before allowing the image to be deployed. At deploy time, instead of redoing activities that were completed in earlier stages, Binary Authorization only needs to verify the attestation. <https://cloud.google.com/binary-authorization/docs/overview>

NEW QUESTION 227

- (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 230

- (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 232

- (Topic 5)

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

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

Answer: BE

Explanation:

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

NEW QUESTION 233

- (Topic 5)

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

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

Answer: A

Explanation:

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

Google Cloud Dedicated Interconnect provides direct physical connections and RFC 1918 communication between your on-premises network and Google's network. Dedicated Interconnect enables you to transfer large amounts of data between networks, which can be more cost effective than purchasing additional bandwidth over the public Internet or using VPN tunnels.

Benefits:

? Traffic between your on-premises network and your VPC network doesn't traverse the public Internet. Traffic traverses a dedicated connection with fewer hops, meaning there are less points of failure where traffic might get dropped or disrupted.

? Your VPC network's internal (RFC 1918) IP addresses are directly accessible from your on-premises network. You don't need to use a NAT device or VPN tunnel to reach internal IP addresses. Currently, you can only reach internal IP addresses over a dedicated connection. To reach Google external IP addresses, you must use a separate connection.

? You can scale your connection to Google based on your needs. Connection capacity is delivered over one or more 10 Gbps Ethernet connections, with a maximum of eight connections (80 Gbps total per interconnect).

? The cost of egress traffic from your VPC network to your on-premises network is reduced. A dedicated connection is generally the least expensive method if you have a high-volume of traffic to and from Google's network.

References: <https://cloud.google.com/interconnect/docs/details/dedicated>

NEW QUESTION 234

- (Topic 5)

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

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

Answer: C

NEW QUESTION 239

- (Topic 5)

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

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

Answer: B

Explanation:

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

NEW QUESTION 240

- (Topic 5)

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

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

Answer: AE

NEW QUESTION 244

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 247

- (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 252

- (Topic 5)

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

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

Answer: A

Explanation:

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

? <https://cloud.google.com/data-transfer-appliance/docs/overview>

? <https://cloud.google.com/blog/products/storage-data-transfer/introducing-storage-transfer-service-for-on-premises-data>

NEW QUESTION 254

- (Topic 5)

You are monitoring Google Kubernetes Engine (GKE) clusters in a Cloud Monitoring workspace. As a Site Reliability Engineer (SRE), you need to triage incidents quickly. What should you do?

- A. Navigate the predefined dashboards in the Cloud Monitoring workspace, and then add metrics and create alert policies.
- B. Navigate the predefined dashboards in the Cloud Monitoring workspace, create custom metrics, and install alerting software on a Compute Engine instance.
- C. Write a shell script that gathers metrics from GKE nodes, publish these metrics to a Pub/Sub topic, export the data to BigQuery, and make a Data Studio dashboard.
- D. Create a custom dashboard in the Cloud Monitoring workspace for each incident, and then add metrics and create alert policies.

Answer: A

Explanation:

<https://cloud.google.com/stackdriver/docs/solutions/gke/legacy-stackdriver/monitoring>

NEW QUESTION 258

- (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 260

- (Topic 5)

A production database virtual machine on Google Compute Engine has an ext4-formatted persistent disk for data files. The database is about to run out of storage space. How can you remediate the problem with the least amount of downtime?

- A. In the Cloud Platform Console, increase the size of the persistent disk and use the `resize2fs` command in Linux.
- B. Shut down the virtual machine, use the Cloud Platform Console to increase the persistent disk size, then restart the virtual machine.
- C. In the Cloud Platform Console, increase the size of the persistent disk and verify the new space is ready to use with the `fdisk` command in Linux.
- D. In the Cloud Platform Console, create a new persistent disk attached to the virtual machine, format and mount it, and configure the database service to move the files to the new disk.
- E. In the Cloud Platform Console, create a snapshot of the persistent disk, restore the snapshot to a new larger disk, unmount the old disk, mount the new disk, and restart the database service.

Answer: A

Explanation:

On Linux instances, connect to your instance and manually resize your partitions and file systems to use the additional disk space that you added.

Extend the file system on the disk or the partition to use the added space. If you grew a partition on your disk, specify the partition. If your disk does not have a partition table, specify only the disk ID.

```
sudo resize2fs /dev/[DISK_ID][PARTITION_NUMBER]
```

where [DISK_ID] is the device name and [PARTITION_NUMBER] is the partition number for the device where you are resizing the file system.

References: <https://cloud.google.com/compute/docs/disks/add-persistent-disk>

NEW QUESTION 264

- (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 268

- (Topic 5)

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

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

Answer: C

Explanation:

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

NEW QUESTION 273

- (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 278

- (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 280

- (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 layers
- 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 's size each day until autoscale logic is triggered on all layers
- 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 285

- (Topic 5)

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

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

Answer: D

NEW QUESTION 287

- (Topic 5)

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

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

Answer: C

Explanation:

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

References: <https://cloud.google.com/logging/>

NEW QUESTION 292

- (Topic 5)

Your web application has several VM instances running within a VPC. You want to restrict communications between instances to only the paths and ports you authorize, but you don't want to rely on static IP addresses or subnets because the app can autoscale. How should you restrict communications?

- A. Use separate VPCs to restrict traffic
- B. Use firewall rules based on network tags attached to the compute instances
- C. Use Cloud DNS and only allow connections from authorized hostnames
- D. Use service accounts and configure the web application particular service accounts to have access

Answer: B

NEW QUESTION 294

- (Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 297

- (Topic 5)

You are designing a mobile chat application. You want to ensure people cannot spoof chat messages, by providing a message were sent by a specific user. What should you do

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

Answer: C

NEW QUESTION 300

- (Topic 5)

Your company has successfully migrated to the cloud and wants to analyze their data stream to optimize operations. They do not have any existing code for this analysis, so they are exploring all their options. These options include a mix of batch and stream processing, as they are running some hourly jobs and live-processing some data as it comes in. Which technology should they use for this?

- A. Google Cloud Dataproc
- B. Google Cloud Dataflow
- C. Google Container Engine with Bigtable
- D. Google Compute Engine with Google BigQuery

Answer: B

Explanation:

Dataflow is for processing both the Batch and Stream.

Cloud Dataflow is a fully-managed service for transforming and enriching data in stream (real time) and batch (historical) modes with equal reliability and expressiveness -- no more complex workarounds or compromises needed.

References: <https://cloud.google.com/dataflow/>

NEW QUESTION 301

- (Topic 5)

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

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

Answer: D

Explanation:

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

NEW QUESTION 302

- (Topic 6)

For this question, refer to the Dress4Win case study. Considering the given business requirements, how would you automate the deployment of web and transactional data layers?

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

Answer: A

NEW QUESTION 305

- (Topic 7)

TerraEarth 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-channel>

NEW QUESTION 306

- (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 307

- (Topic 7)

For this question, refer to the TerramEarth case study. Considering the technical requirements, how should you reduce the unplanned vehicle downtime in GCP?

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

Answer: A

NEW QUESTION 311

- (Topic 7)

You are migrating a Linux-based application from your private data center to Google Cloud. The TerramEarth 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 312

- (Topic 7)

For this question, refer to the TerramEarth case study. You are asked to design a new architecture for the ingestion of the data of the 200,000 vehicles that are connected to a cellular network. You want to follow Google-recommended practices.

Considering the technical requirements, which components should you use for the ingestion of the data?

- A. Google Kubernetes Engine with an SSL Ingress
- B. Cloud IoT Core with public/private key pairs
- C. Compute Engine with project-wide SSH keys
- D. Compute Engine with specific SSH keys

Answer: A

Explanation:

<https://cloud.google.com/solutions/iot-overview> <https://cloud.google.com/iot/quotas>

NEW QUESTION 315

- (Topic 7)

For this question, refer to the TerramEarth case study. You need to implement a reliable, scalable GCP solution for the data warehouse for your company,

TerramEarth. Considering the TerramEarth business and technical requirements, what should you do?

- A. Replace the existing data warehouse with BigQuer
- B. Use table partitioning.
- C. Replace the existing data warehouse with a Compute Engine instance with 96 CPUs.
- D. Replace the existing data warehouse with BigQuer
- E. Use federated data sources.
- F. Replace the existing data warehouse with a Compute Engine instance with 96 CPU
- G. Add an additional Compute Engine pre-emptible instance with 32 CPUs.

Answer: C

Explanation:

https://cloud.google.com/solutions/bigquery-data-warehouse#external_sources <https://cloud.google.com/solutions/bigquery-data-warehouse>

NEW QUESTION 317

- (Topic 8)

For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to design their solution for the future in order to take advantage of cloud and technology improvements as they become available. Which two steps should they take? (Choose two.)

- A. Store as much analytics and game activity data as financially feasible today so it can be used to train machine learning models to predict user behavior in the future.
- B. Begin packaging their game backend artifacts in container images and running them on Kubernetes Engine to improve the availability to scale up or down based on game activity.
- C. Set up a CI/CD pipeline using Jenkins and Spinnaker to automate canary deployments and improve development velocity.
- D. Adopt a schema versioning tool to reduce downtime when adding new game features that require storing additional player data in the database.
- E. Implement a weekly rolling maintenance process for the Linux virtual machines so they can apply critical kernel patches and package updates and reduce the risk of 0-day vulnerabilities.

Answer: BC

NEW QUESTION 318

- (Topic 8)

For this question, refer to the Mountkirk Games case study. Which managed storage option meets Mountkirk's technical requirement for storing game activity in a time series database service?

- A. Cloud Bigtable
- B. Cloud Spanner
- C. BigQuery
- D. Cloud Datastore

Answer: A

Explanation:

<https://cloud.google.com/blog/products/databases/getting-started-with-time-series-trend-predictions-using-gcp>

NEW QUESTION 321

- (Topic 8)

For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to migrate from their current analytics and statistics reporting model to one that meets their technical requirements on Google Cloud Platform.

Which two steps should be part of their migration plan? (Choose two.)

- A. Evaluate the impact of migrating their current batch ETL code to Cloud Dataflow.
- B. Write a schema migration plan to denormalize data for better performance in BigQuery.
- C. Draw an architecture diagram that shows how to move from a single MySQL database to a MySQL cluster.
- D. Load 10 TB of analytics data from a previous game into a Cloud SQL instance, and run test queries against the full dataset to confirm that they complete successfully.
- E. Integrate Cloud Armor to defend against possible SQL injection attacks in analytics files uploaded to Cloud Storage.

Answer: AB

Explanation:

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

NEW QUESTION 324

- (Topic 8)

For this question, refer to the Mountkirk Games case study. Mountkirk Games wants you to design a way to test the analytics platform's resilience to changes in mobile network latency. What should you do?

- A. Deploy failure injection software to the game analytics platform that can inject additional latency to mobile client analytics traffic.
- B. Build a test client that can be run from a mobile phone emulator on a Compute Engine virtual machine, and run multiple copies in Google Cloud Platform regions all over the world to generate realistic traffic.
- C. Add the ability to introduce a random amount of delay before beginning to process analytics files uploaded from mobile devices.
- D. Create an opt-in beta of the game that runs on players' mobile devices and collects response times from analytics endpoints running in Google Cloud Platform regions all over the world.

Answer: D

NEW QUESTION 329

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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