

Google.Professional-Cloud-Architect.v2023-10-16.q124

Exam Code:	Professional-Cloud-Architect
Exam Name:	Google Certified Professional - Cloud Architect (GCP)
Certification Provider:	Google
Free Question Number:	124
Version:	v2023-10-16
# of views:	1846
# of Questions views:	1240
https://www.exam-tests.com/Professional-Cloud-Architect-exam/Google.Professional-Cloud-Architect.v2023-10-16.q124.html	

NEW QUESTION: 1

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: ([SHOW ANSWER](#))

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

NEW QUESTION: 2

You are working in a highly secured environment where public Internet access from the Compute Engine VMs is not allowed. You do not yet have a VPN connection to access an on-premises file server. You need to install specific software on a Compute Engine instance. How should you install the software?

- A. Upload the required installation files to Cloud Storage. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gsutil.
- B. Upload the required installation files to Cloud Storage and use firewall rules to block all traffic except the IP address range for Cloud Storage. Download the files to the VM using gsutil.
- C. Upload the required installation files to Cloud Source Repositories. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gcloud.

D. Upload the required installation files to Cloud Source Repositories and use firewall rules to block all traffic except the IP address range for Cloud Source Repositories. Download the files to the VM using gsutil.

Answer: A ([LEAVE A REPLY](#))

<https://cloud.google.com/vpc/docs/private-access-options#pga-supported>

NEW QUESTION: 3

Your company has just recently activated Cloud Identity to manage users. The Google Cloud Organization has been configured as web. The security team needs to secure protects that will be part of the Organization. They want to prohibit IAM users outside the domain from gaining permissions from now on. What should they do?

- A.** Configure Cloud Scheduler to trigger a Cloud Function every hour that removes all users that don't belong to the Cloud identity domain from all projects.
- B.** Configure an organization policy to restrict identities by domain
- C.** Configure an organization policy to block creation of service accounts

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 4

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

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

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 5

For this question, refer to the TerramEarth case study.

TerramEarth plans to connect all 20 million vehicles in the field to the cloud. This increases the volume to 20 million 600 byte records a second for 40 TB an hour. How should you design the data ingestion?

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

Answer: B ([LEAVE A REPLY](#))

<https://cloud.google.com/solutions/data-lifecycle-cloud-platform>

<https://cloud.google.com/solutions/designing-connected-vehicle-platform>

NEW QUESTION: 6

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: ([SHOW ANSWER](#))

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

NEW QUESTION: 7

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. Deploy a Cloud SQL server to replace MySQL. Deploy Jenkins using Cloud Deployment Manager.
- B. Deploy Nginx and Tomcat using Cloud Launcher. Deploy a MySQL server using Cloud Launcher. Deploy Jenkins to Compute Engine using Cloud Deployment Manager scripts.
- C. Migrate Nginx and Tomcat to App Engine. Deploy a Cloud Datastore server to replace the MySQL server in a high-availability configuration. Deploy Jenkins to Compute Engine using Cloud Launcher.
- D. Migrate Nginx and Tomcat to App Engine. Deploy a MySQL server using Cloud Launcher. Deploy Jenkins to Compute Engine using Cloud Launcher.

Answer: A ([LEAVE A REPLY](#))

Topic 7, TerramEarth Case 2

Company Overview

TerramEarth manufactures heavy equipment for the mining and agricultural industries. About 80% of their business is from mining and 20% from agriculture. They currently have over 500 dealers and service centers in 100 countries. Their mission is to build products that make their customers more productive.

Solution Concept

There are 20 million TerramEarth vehicles in operation that collect 120 fields of data per second. Data is stored locally on the vehicle and can be accessed for analysis when a vehicle is serviced. The data is downloaded via a maintenance port. This same port can be used to adjust operational parameters, allowing the vehicles to be upgraded in the field with new computing modules.

Approximately 200,000 vehicles are connected to a cellular network, allowing TerramEarth to collect data directly. At a rate of 120 fields of data per second with 22 hours of operation per day, TerramEarth collects a total of about 9 TB/day from these connected vehicles.

Existing Technical Environment

TerramEarth's existing architecture is composed of Linux and Windows-based systems that reside in a single U.S. west coast based data center. These systems gzip CSV files from the field and upload via FTP, and place the data in their data warehouse. Because this process takes time, aggregated reports are based on data that is 3 weeks old.

With this data, TerramEarth has been able to preemptively stock replacement parts and reduce unplanned downtime of their vehicles by 60%. However, because the data is stale, some customers are without their vehicles for up to 4 weeks while they wait for replacement parts.

Business Requirements

Decrease unplanned vehicle downtime to less than 1 week.

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.

Technical Requirements

Expand beyond a single datacenter to decrease latency to the American Midwest and east coast.

Create a backup strategy.

Increase security of data transfer from equipment to the datacenter.

Improve data in the data warehouse.

Use customer and equipment data to anticipate customer needs.

Application 1: Data ingest

A custom Python application reads uploaded datafiles from a single server, writes to the data warehouse.

Compute:

Windows Server 2008 R2

- 16 CPUs
- 128 GB of RAM
- 10 TB local HDD storage

Application 2: Reporting

An off the shelf application that business analysts use to run a daily report to see what equipment needs repair. Only 2 analysts of a team of 10 (5 west coast, 5 east coast) can connect to the reporting application at a time.

Compute:

Off the shelf application. License tied to number of physical CPUs

- Windows Server 2008 R2
- 16 CPUs
- 32 GB of RAM
- 500 GB HDD

Data warehouse:

A single PostgreSQL server

- RedHat Linux
- 64 CPUs
- 128 GB of RAM
- 4x 6TB HDD in RAID 0

Executive Statement

Our competitive advantage has always been in the manufacturing process, with our ability to build better vehicles for lower cost than our competitors. However, new products with different approaches are constantly being developed, and I'm concerned that we lack the skills to undergo the next wave of transformations in our industry. My goals are to build our skills while addressing immediate market needs through incremental innovations.

NEW QUESTION: 8

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. Protect the bucket with a password that expires after 24 hours.
- B.** Have users upload the images to Cloud Storage using a signed URL that expires after 24 hours.
- C.** Create an App Engine web application where users can upload images. Configure App Engine to disable the application after 24 hours. Authenticate users via Cloud Identity.
- D.** Create an App Engine web application where users can upload images for the next 24 hours. Authenticate users via Cloud Identity.

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 9

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

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

Answer: B,C ([LEAVE A REPLY](#))

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

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

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

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

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

NEW QUESTION: 10

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 (LEAVE A REPLY)

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

NEW QUESTION: 11

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

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

Answer: (SHOW ANSWER)

NEW QUESTION: 12

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 ([LEAVE A REPLY](#))

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.

NEW QUESTION: 13

For this question, refer to the EHR Healthcare case study. In the past, configuration errors put public IP addresses on backend servers that should not have been accessible from the Internet. You need to ensure that no one can put external IP addresses on backend Compute Engine instances and that external IP addresses can only be configured on frontend Compute Engine instances. What should you do?

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

Answer: A ([LEAVE A REPLY](#))

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

NEW QUESTION: 14

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 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
- 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 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
- 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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 15

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

- A.** Upload your files into Cloud Storage. Use Cloud Datalab to explore and clean your data.
- B.** Upload your files into Cloud Storage. Use Cloud Dataprep to explore and clean your data.
- C.** Connect Cloud Datalab to your on-premises systems. Use Cloud Datalab to explore and clean your data.
- D.** Connect Cloud Dataprep to your on-premises systems. Use Cloud Dataprep to explore and clean your

Answer: **B** ([LEAVE A REPLY](#))

data.

Explanation:

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

NEW QUESTION: 16

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

- A.** 1. Create a VPC Service Controls perimeter that includes the projects with the buckets.
2. Create an access level with the CIDR of the office network.
- B.** 1. Create a firewall rule for all instances in the Virtual Private Cloud (VPC) network for source range.
2. Use the Classless Inter-domain Routing (CIDR) of the office network.

- C.** 1. Create a Cloud Function to remove IAM permissions from the buckets, and another Cloud Function to add IAM permissions to the buckets.
2. Schedule the Cloud Functions with Cloud Scheduler to add permissions at the start of business and remove permissions at the end of business.
- D.** 1. Create a Cloud VPN to the office network.
2. Configure Private Google Access for on-premises hosts.

Answer: A (LEAVE A REPLY)

For all Google Cloud services secured with VPC Service Controls, you can ensure that: Resources within a perimeter are accessed only from clients within authorized VPC networks using Private Google Access with either Google Cloud or on-premises.

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

<https://cloud.google.com/vpc-service-controls/docs/overview>. You create a service control across your VPC and any cloud bucket or any project resource to restrict access. Anything outside of it can't access the resources within service control perimeter

Valid Professional-Cloud-Architect Dumps shared by BraindumpsPass.com for Helping Passing Professional-Cloud-Architect Exam! BraindumpsPass.com now offer the **newest Professional-Cloud-Architect exam dumps**, the BraindumpsPass.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com Professional-Cloud-Architect dumps with Test Engine here: <https://www.braindumpsPASS.com/Google/Professional-Cloud-Architect-practice-exam-dumps.html> (378 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 17

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

What should you do?

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

Answer: (SHOW ANSWER)

Topic 8, Mountkrik Games Case 2

Company Overview

Mountkirk Games makes online, session-based, multiplayer games for mobile platforms. They build all of their games using some server-side integration. Historically, they have used cloud providers to lease physical servers.

Due to the unexpected popularity of some of their games, they have had problems scaling their global audience, application servers, MySQL databases, and analytics tools.

Their current model is to write game statistics to files and send them through an ETL tool that loads them into a centralized MySQL database for reporting.

Solution Concept

Mountkirk Games is building a new game, which they expect to be very popular. They plan to deploy the game's backend on Google Compute Engine so they can capture streaming metrics, run intensive analytics, and take advantage of its autoscaling server environment and integrate with a managed NoSQL database.

Business Requirements

Increase to a global footprint.

Improve uptime - downtime is loss of players.

Increase efficiency of the cloud resources we use.

Reduce latency to all customers.

Technical Requirements

Requirements for Game Backend Platform

Dynamically scale up or down based on game activity.

Connect to a transactional database service to manage user profiles and game state.

Store game activity in a timeseries database service for future analysis.

As the system scales, ensure that data is not lost due to processing backlogs.

Run hardened Linux distro.

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 queries to access at least 10 TB of historical data

Process files that are regularly uploaded by users' mobile devices


Executive Statement

Our last successful game did not scale well with our previous cloud provider, resulting in lower user adoption and affecting the game's reputation. Our investors want more key performance indicators (KPIs) to evaluate the speed and stability of the game, as well as other metrics that provide deeper insight into usage patterns so we can adapt the game to target users. Additionally, our current technology stack cannot provide the scale we need, so we want to replace MySQL and move to an environment that provides autoscaling, low latency load balancing, and frees us up from managing physical servers.

NEW QUESTION: 18

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)

Answer: A,C (LEAVE A REPLY)



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

C. Compute Engine

D. GKE On-Prem

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 20

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

A. Use the `gsutil mv` command to move the data

B. Use the Storage Transfer Service to move the data

C. Download the data to a Transfer Appliance and ship it to Google

D. Download the data to the on-premises data center and upload it to the Cloud Storage bucket

Answer: B ([LEAVE A REPLY](#))

<https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#transfer-options>

<https://cloud.google.com/storage-transfer-service>

NEW QUESTION: 21

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

A. Use `kubectl set image deployment/echo-deployment <new-image>`

B. Use the rolling update functionality of the Instance Group behind the Kubernetes cluster

C. Update the deployment yaml file with the new container image. Use `kubectl delete deployment/ echo-deployment` and `kubectl create -f <yaml-file>`

D. Update the service yaml file with the new container image. Use `kubectl delete service/echoservice` and `kubectl create -f <yaml-file>`

Answer: A ([LEAVE A REPLY](#))

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

NEW QUESTION: 22

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

A. Configure HashiCorp Vault on Compute Engine, and use customer managed encryption keys and Cloud Key Management Service (Cloud KMS) to manage the encryption keys. Configure these Secrets to be used by the application platform.

- B.** Configure Kubernetes Secrets to store the secret, enable Application-Layer Secrets Encryption, and use Cloud Key Management Service (Cloud KMS) to manage the encryption keys. Configure these Secrets to be used by the application platform.
- C.** Use Kubernetes Secrets, which are obfuscated by default. Configure these Secrets to be used by the application platform.
- D.** Configure Workload Identity and service accounts to be used by the application platform.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 23

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: A,B ([LEAVE A REPLY](#))

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

NEW QUESTION: 24

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

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

Answer: C ([LEAVE A REPLY](#))

<https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy>

"Roles are always inherited, and there is no way to explicitly remove a permission for a lower-level resource that is granted at a higher level in the resource hierarchy. Given the above example, even if you were to remove the Project Editor role from Bob on the "Test GCP Project",

he would still inherit that role from the "Dept Y" folder, so he would still have the permissions for that role on "Test GCP Project".

NEW QUESTION: 25

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

- A.** Create a scheduled job in Cloud Run to invoke a container every minute. 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.
- B.** Create a cron job on a Compute Engine VM that runs every minute. 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.
- C.** 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.
- D.** 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: ([SHOW ANSWER](#))

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

NEW QUESTION: 26

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

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

Answer: **A** ([LEAVE A REPLY](#))

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

NEW QUESTION: 27

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 registered as a xinetd service in Linux and configure a Stackdriver endpoint check to call the service.
- B.** 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.
- C.** Create a shutdown script named k99.shutdown in the /etc/rc.6.d/ directory.
- 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: B ([LEAVE A REPLY](#))

NEW QUESTION: 28

As part of implementing their disaster recovery plan, your company is trying to replicate their production MySQL database from their private data center to their GCP project using a Google Cloud VPN connection.

They are experiencing latency issues and a small amount of packet loss that is disrupting the replication. What should they do?

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

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 29

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

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 30

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

A. Perform the following:

- 1) Create a managed instance group with f1-micro type machines.
- 2) Use a startup script to clone the repository, check out the production branch, install the dependencies, and start the Python app.
- 3) Restart the instances to automatically deploy new production releases.

B. Perform the following:

- 1) Create a managed instance group with n1-standard-1 type machines.
- 2) Build a Compute Engine image from the production branch that contains all of the dependencies and automatically starts the Python app.
- 3) Rebuild the Compute Engine image, and update the instance template to deploy new production releases.

C. Perform the following:

- 1) Create a Kubernetes Engine cluster with n1-standard-1 type machines.
- 2) Build a Docker image from the production branch with all of the dependencies, and tag it with the
- 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 will all of the dependencies, and tag it with "latest".
- 3) Create a Kubernetes Deployment in the default namespace with the imagePullPolicy set to "Always".Restart the pods to automatically deploy new production releases.

Answer: D ([LEAVE A REPLY](#))

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

NEW QUESTION: 31

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

What should you do?

A. Deploy intrusion detection software to your virtual machines to detect and log unauthorized access.

B. Configure a red replica for your Cloud SQL instance in a different zone than the master, and then manually trigger a failover while monitoring KPIs for our REST API.

- C. Engage with a security company to run web scrapes that look your users' authentication data on malicious websites and notify you if any is found.
- D. Schedule a disaster simulation exercise during which you can shut off all VMs in a zone to see how your application behaves.

Answer: ([SHOW ANSWER](#))

Valid Professional-Cloud-Architect Dumps shared by BraindumpsPass.com for Helping Passing Professional-Cloud-Architect Exam! BraindumpsPass.com now offer the **newest Professional-Cloud-Architect exam dumps**, the BraindumpsPass.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com Professional-Cloud-Architect dumps with Test Engine here: <https://www.braindumpsPass.com/Google/Professional-Cloud-Architect-practice-exam-dumps.html> (378 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 32

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

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 33

For this question, refer to the Mountkirk Games case study. You are in charge of the new Game Backend Platform architecture. The game communicates with the backend over a REST API.

You want to follow Google-recommended practices. How should you design the backend?

- A. Create an instance template for the backend. For every region, deploy it on a multi-zone managed instance group. Use an L4 load balancer.
- B. Create an instance template for the backend. For every region, deploy it on a single-zone managed instance group. Use an L4 load balancer.
- C. Create an instance template for the backend. For every region, deploy it on a multi-zone managed instance group. Use an L7 load balancer.
- D. Create an instance template for the backend. For every region, deploy it on a single-zone managed instance group. Use an L7 load balancer.

Answer: C ([LEAVE A REPLY](#))

https://cloud.google.com/solutions/gaming/cloud-game-infrastructure#dedicated_game_server

NEW QUESTION: 34

A development team at your company has created a dockerized HTTPS web application. You need to deploy the application on Google Kubernetes Engine (GKE) and make sure that the application scales automatically.

How should you deploy to GKE?

- A.** Use the Horizontal Pod Autoscaler and enable cluster autoscaling. Use an Ingress resource to loadbalance the HTTPS traffic.
- B.** Use the Horizontal Pod Autoscaler and enable cluster autoscaling on the Kubernetes cluster. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.
- C.** Enable autoscaling on the Compute Engine instance group. Use an Ingress resource to load balance the HTTPS traffic.
- D.** Enable autoscaling on the Compute Engine instance group. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.

Answer: B ([LEAVE A REPLY](#))

<https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer>

<https://cloud.google.com/kubernetes-engine/docs/concepts/network-overview#ext-lb>

NEW QUESTION: 35

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: (SHOW ANSWER)

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

NEW QUESTION: 36

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 ([LEAVE A REPLY](#))

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.

Incorrect Answers:

A: HTTP server push, also known as HTTP streaming, is a client-server communication pattern that sends information from an HTTP server to a client asynchronously, without a client request. A server push architecture is especially effective for highly interactive web or mobile applications, where one or more clients need to receive continuous information from the server.

NEW QUESTION: 37

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

Answer: (SHOW ANSWER)

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

NEW QUESTION: 38

You are working at an institution that processes medical data.

a. You are migrating several workloads onto Google Cloud. Company policies require all workloads to run on physically separated hardware, and workloads from different clients must also be separated. You created a sole-tenant node group and added a node for each client. You need to deploy the workloads on these dedicated hosts. What should you do?

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

Answer: C (LEAVE A REPLY)

https://cloud.google.com/compute/docs/nodes/provisioning-sole-tenant-vms#provision_a_sole-tenant_vm

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

NEW QUESTION: 39

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

- A.** Create a snapshot schedule for the disk containing the application data. Whenever a zonal outage occurs, use the latest snapshot to restore the disk in the same zone.
- B.** Configure the Compute Engine instances with an instance template for the application, and use a regional persistent disk for the application data. Whenever a zonal outage occurs, use the instance template to spin up the application in another zone in the same region. Use the regional persistent disk for the application data.
- C.** Create a snapshot schedule for the disk containing the application data. Whenever a zonal outage occurs, use the latest snapshot to restore the disk in another zone within the same region.

D. Configure the Compute Engine instances with an instance template for the application, and use a regional persistent disk for the application data. Whenever a zonal outage occurs, use the instance template to spin up the application in another region. Use the regional persistent disk for the application data,

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 40

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

- A.** Allocate budget for team training. Set a deadline for the new GCP project.
- B.** Allocate budget for team training. Create a roadmap for your team to achieve Google Cloud certification based on job role.
- C.** Allocate budget to hire skilled external consultants. Set a deadline for the new GCP project.
- D.** Allocate budget to hire skilled external consultants. Create a roadmap for your team to achieve Google Cloud certification based on job role.

Answer: **B** ([LEAVE A REPLY](#))

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

NEW QUESTION: 41

For this question, refer to the JencoMart case study.

The JencoMart security team requires that all Google Cloud Platform infrastructure is deployed using a least privilege model with separation of duties for administration between production and development resources. What Google domain and project structure should you recommend?

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

Answer: **D** ([LEAVE A REPLY](#))

Note: The principle of least privilege and separation of duties are concepts that, although semantically different, are intrinsically related from the standpoint of security. The intent behind both is to prevent people from having higher privilege levels than they actually need Principle of Least Privilege: Users should only have the least amount of privileges required to perform their

job and no more. This reduces authorization exploitation by limiting access to resources such as targets, jobs, or monitoring templates for which they are not authorized.

Separation of Duties: Beyond limiting user privilege level, you also limit user duties, or the specific jobs they can perform. No user should be given responsibility for more than one related function. This limits the ability of a user to perform a malicious action and then cover up that action.

NEW QUESTION: 42

Your company has an application deployed on Anthos clusters (formerly Anthos GKE) that is running multiple microservices. The cluster has both Anthos Service Mesh and Anthos Config Management configured. End users inform you that the application is responding very slowly. You want to identify the microservice that is causing the delay. What should you do?

- A.** Use the Service Mesh visualization in the Cloud Console to inspect the telemetry between the microservices.
- B.** Use Anthos Config Management to create a ClusterSelector selecting the relevant cluster. On the Google Cloud Console page for Google Kubernetes Engine, view the Workloads and filter on the cluster. Inspect the configurations of the filtered workloads.
- C.** Use Anthos Config Management to create a namespaceSelector selecting the relevant cluster namespace. On the Google Cloud Console page for Google Kubernetes Engine, visit the workloads and filter on the namespace. Inspect the configurations of the filtered workloads.
- D.** Reinstall istio using the default istio profile in order to collect request latency. Evaluate the telemetry between the microservices in the Cloud Console.

Answer: A ([LEAVE A REPLY](#))

The Anthos Service Mesh pages in the Google Cloud Console provide both summary and in-depth metrics, charts, and graphs that enable you to observe service behavior. You can monitor the overall health of your services, or drill down on a specific service to set a service level objective (SLO) or troubleshoot an issue. <https://cloud.google.com/service-mesh/docs/observability/explore-dashboard>
<https://cloud.google.com/anthos/service-mesh>

NEW QUESTION: 43

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

```
FROM ubuntu:16.04
```

```
COPY . /src
```

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

```
RUN pip install -r requirements.txt
```

You want to optimize this Dockerfile for faster deployment times without adversely affecting the app's functionality.

Which two actions should you take? Choose 2 answers.

- A.** Remove Python after running pip.

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

Answer: C,E (LEAVE A REPLY)

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

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

Reference:

<https://www.alpinelinux.org/about/>

NEW QUESTION: 44

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

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

Answer: B (LEAVE A REPLY)

Current legislation requires you to delete "SUCH" information upon request of the subject. " So from that point of view the question is not to delete the entire user records but specific data related to personal health data. With DLP you can use InfoTypes and InfoType detectors to specifically scan for those entries and how to act upon them (link

<https://cloud.google.com/dlp/docs/concepts-infotypes>)

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

NEW QUESTION: 45

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: ([SHOW ANSWER](#))

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

NEW QUESTION: 46

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 Cloud Dataproc Hive as the data warehouse. Directly stream data into prtitioned Hive tables. Use Pig scripts to analyze data.
- B.** Use BigQuery as the data warehouse. Connect all vehicles to the network and stream data into BigQuery using Cloud Pub/Sub and Cloud Dataflow. Use Google Data Studio for analysis and reporting.
- C.** Use BigQuery as the data warehouse. Connect all vehicles to the network and upload gzip files to a Multi-Regional Cloud Storage bucket using gcloud. Use Google Data Studio for analysis and reporting.
- D.** Use Cloud Dataproc Hive as the data warehouse. Upload gzip files to a MultiRegional Cloud Storage bucket. Upload this data into BigQuery using gcloud. Use Google data Studio for analysis and reporting.

Answer: **B** ([LEAVE A REPLY](#))

Valid Professional-Cloud-Architect Dumps shared by BraindumpsPass.com for Helping Passing Professional-Cloud-Architect Exam! BraindumpsPass.com now offer the **newest Professional-Cloud-Architect exam dumps**, the BraindumpsPass.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com Professional-Cloud-Architect dumps with Test Engine here: <https://www.braindumpsPass.com/Google/Professional-Cloud-Architect-practice-exam-dumps.html> (378 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 47

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. The generated container can then be deployed into a Kubernetes cluster.
- C.** 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.
- D.** The process can be automated with Migrate for Compute Engine.

Answer: C (LEAVE A REPLY)

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: 48

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

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

Answer: B (LEAVE A REPLY)

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

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

NEW QUESTION: 49

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.

Otherwise, save files to existing bucket

B. * 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. Otherwise, save the single archive file to existing bucket.

C. * Compress individual files.

* Name files with serverName-EventSequence.

* Save files to one bucket

* Set custom metadata headers for each object after saving.

D. * Append metadata to file body.

* Compress individual files.

* Name files with a random prefix pattern.

* Save files to one bucket

Answer: ([SHOW ANSWER](#))

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: 50

Your applications will be writing their logs to BigQuery for analysis. Each application should have its own table.

Any logs older than 45 days should be removed. You want to optimize storage and follow Google recommended practices. What should you do?

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

B. Make the tables time-partitioned, and configure the partition expiration at 45 days

C. Rely on BigQuery's default behavior to prune application logs older than 45 days

D. Create a script that uses the BigQuery command line tool (bq) to remove records older than 45 days

Answer: **B** ([LEAVE A REPLY](#))

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

NEW QUESTION: 51

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

What should you do?

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

Answer: A ([LEAVE A REPLY](#))

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

NEW QUESTION: 52

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 53

Your customer support tool logs all email and chat conversations to Cloud Bigtable for retention and analysis.

What is the recommended approach for sanitizing this data of personally identifiable information or payment card information before initial storage?

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

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 54

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. Enable the Compute Engine API and then enable logging on the firewall rules that match the traffic you want to save

- B.** Set up a filter in Cloud Logging and a Cloud Storage bucket as an export target for the logs you want to save
- C.** 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
- D.** Set up a filter in Cloud Logging and a topic in Pub/Sub to publish the logs

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 55

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 ([LEAVE A REPLY](#))

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

NEW QUESTION: 56

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

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

Answer: ([SHOW ANSWER](#))

<https://cloud.google.com/security/gdpr/?tab=tab4>

NEW QUESTION: 57

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.** 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.
- B.** 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.
- C.** 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.
- D.** Adopt a schema versioning tool to reduce downtime when adding new game features that require storing additional player data in the database.
- E.** Set up a CI/CD pipeline using Jenkins and Spinnaker to automate canary deployments and improve development velocity.

Answer: A,E ([LEAVE A REPLY](#))

NEW QUESTION: 58

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: ([SHOW ANSWER](#))

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

NEW QUESTION: 59

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. Structure all projects for each department under the respective folders.

Answer: A ([LEAVE A REPLY](#))

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). Reference:

<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: 60

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

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 61

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 instance in vpc-gcloud

Answer: C (LEAVE A REPLY)

```
gcloud compute start-iap-tunnel INSTANCE_NAME_IN_VPC_B 22 \  
--local-host-port=localhost:22
```

1. Create an additional instance in vpc-a
2. Create an additional instance n vpc-b
3. Instal OpenVPN in newly created instances
4. Configure a VPN tunnel between vpc-a and vpc-b with the help of OpenVPN

Valid Professional-Cloud-Architect Dumps shared by BraindumpsPass.com for Helping Passing Professional-Cloud-Architect Exam! BraindumpsPass.com now offer the **newest Professional-Cloud-Architect exam dumps**, the BraindumpsPass.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com Professional-Cloud-Architect dumps with Test Engine here: <https://www.braindumpsPASS.com/Google/Professional-Cloud-Architect-practice-exam-dumps.html> (378 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 62

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

A. 1. From the dataset where you have the source data, create views of tables that you want to share, excluding PII.

2. Assign an appropriate project-level IAM role to the members of the data science team.

3. Assign access controls to the dataset that contains the view.

B. 1. From the dataset where you have the source data, create materialized views of tables that you want to share, excluding PII.

2. Assign an appropriate project-level IAM role to the members of the data science team. 3. Assign access controls to the dataset that contains the view.

C. 1. Create a dataset for the data science team.

2. Create views of tables that you want to share, excluding PII.

3. Assign an appropriate project-level IAM role to the members of the data science team.

4. Assign access controls to the dataset that contains the view.

5. Authorize the view to access the source dataset.

D. 1. Create a dataset for the data science team.

2. Create materialized views of tables that you want to share, excluding PII
3. Assign an appropriate project-level IAM role to the members of the data science team
4. Assign access controls to the dataset that contains the view
5. Authorize the view to access the source dataset

Answer: C ([LEAVE A REPLY](#))

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

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

<https://cloud.google.com/bigquery/docs/authorized-datasets>

NEW QUESTION: 63

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

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

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

Answer: ([SHOW ANSWER](#)**)**

From scenario: Requirements for Game Backend Platform

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

NEW QUESTION: 64

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

What should you do?

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

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 65

All compute Engine instances in your VPC should be able to connect to an Active Directory server on specific ports. Any other traffic emerging from your instances is not allowed. You want to enforce this using VPC firewall rules.

How should you configure the firewall rules?

- A.** Create an egress rule with priority 1000 to deny all traffic for all instances. Create another egress rule with priority 100 to allow the Active Directory traffic for all instances.
- B.** Create an egress rule with priority 100 to deny all traffic for all instances. Create another egress rule with priority 1000 to allow the Active Directory traffic for all instances.
- C.** Create an egress rule with priority 1000 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 100 to block all traffic for all instances.
- D.** Create an egress rule with priority 100 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 1000 to block all traffic for all instances.

Answer: B ([LEAVE A REPLY](#))

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

NEW QUESTION: 66

For this question, refer to the Dress4Win case study.

Dress4Win has configured a new uptime check with Google Stackdriver for several of their legacy services. The Stackdriver dashboard is not reporting the services as healthy. What should they do?

- A.** Configure their legacy web servers to allow requests that contain user-Agent HTTP header when the value matches GoogleStackdriverMonitoring- UptimeChecks

(<https://cloud.google.com/monitoring>)

- B.** Configure their load balancer to pass through the User-Agent HTTP header when the value matches GoogleStackdriverMonitoring-UptimeChecks (<https://cloud.google.com/monitoring>)
- C.** Install the Stackdriver agent on all of the legacy web servers.
- D.** In the Cloud Platform Console download the list of the uptime servers' IP addresses and create an inbound firewall rule

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 67

You need to upload files from your on-premises environment to Cloud Storage. You want the files to be encrypted on Cloud Storage using customer-supplied encryption keys. What should you do?

- A.** Supply the encryption key in a .boto configuration file. Use gsutil to upload the files.
- B.** Supply the encryption key using gcloud config. Use gsutil to upload the files to that bucket.
- C.** Use gsutil to upload the files, and use the flag --encryption-key to supply the encryption key.
- D.** Use gsutil to create a bucket, and use the flag --encryption-key to supply the encryption key. Use gsutil to upload the files to that bucket.

Answer: A ([LEAVE A REPLY](#))

<https://cloud.google.com/storage/docs/encryption/customer-supplied-keys#gsutil>

NEW QUESTION: 68

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

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

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 69

For this question, refer to the JencoMart case study.

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

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

Answer: D ([LEAVE A REPLY](#))

<https://cloud.google.com/datastore/docs/concepts/overview>

Common workloads for Google Cloud Datastore:

User profiles

Product catalogs

Game state

Reference:

<https://cloud.google.com/datastore/docs/concepts/overview>

NEW QUESTION: 70

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 ([LEAVE A REPLY](#))

<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.

NEW QUESTION: 71

You need to develop procedures to verify resilience of disaster recovery for remote recovery using GCP. Your production environment is hosted on-premises. You need to establish a secure, redundant connection between your on premises network and the GCP network.

What should you do?

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

Answer: B ([LEAVE A REPLY](#))

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

NEW QUESTION: 72

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

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

Answer: A ([LEAVE A REPLY](#))

Create a Compute engine instance with CPU and Memory options similar to your application's current on-premises virtual machine. Install the cloud monitoring agent, and deploy the third party application. Run a load with normal traffic levels on third party application and follow the Rightsizing Recommendations in the Cloud Console

<https://cloud.google.com/migrate/compute-engine/docs/4.9/concepts/planning-a-migration/cloud-instance-rightsizing?hl=en>

NEW QUESTION: 73

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 ([LEAVE A REPLY](#))

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

NEW QUESTION: 74

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

Answer: C ([LEAVE A REPLY](#))

required information.

Explanation:

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

NEW QUESTION: 75

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 ([LEAVE A REPLY](#))

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.

NEW QUESTION: 76

For this question, refer to the JencoMart case study.

JencoMart has built a version of their application on Google Cloud Platform that serves traffic to Asia. You want to measure success against their business and technical goals. Which metrics should you track?

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

Answer: E ([LEAVE A REPLY](#))

Valid Professional-Cloud-Architect Dumps shared by BraindumpsPass.com for Helping Passing Professional-Cloud-Architect Exam! BraindumpsPass.com now offer the **newest Professional-Cloud-Architect exam dumps**, the BraindumpsPass.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com Professional-Cloud-Architect dumps with Test Engine here: <https://www.braindumpsPass.com/Google/Professional-Cloud-Architect-practice-exam-dumps.html> (378 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 77

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

- A. Create an instance with the latest available Debian image. Connect to the instance via SSH, and install and configure the application on the instance. Repeat this process whenever a new Google-managed Debian image becomes available.
- B. Create a Debian-based Compute Engine instance, install and configure the application, and use OS patch management to install available updates.

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

D. Create a Compute Engine instance template using the most recent Debian image. Create an instance from this template, and install and configure the application as part of the startup script. Repeat this process whenever a new Google-managed Debian image becomes available.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 78

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

A. Deploy changes to a small subset of users before rolling out to production.

B. Deploy fewer changes to production.

C. Deploy smaller changes to production.

D. Increase the load on your test and staging environments.

Answer: **D** ([LEAVE A REPLY](#))

NEW QUESTION: 79

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 enable google stack driver debugger on the application code to show errors in the code

B. They should add canary tests so developers can measure how much of an impact the new release causes to latency

C. They should add additional unit tests and production scale load tests on their cloud staging environment.

D. They should run the end to end tests in the cloud staging environment to determine if the code is working as intended.

Answer: **A** ([LEAVE A REPLY](#))

NEW QUESTION: 80

You are implementing Firestore for Mountkirk Games. Mountkirk Games wants to give a new game programmatic access to a legacy game's Firestore database. Access should be as restricted as possible. What should you do?

A. Create a service account (SA) in the legacy game's Google Cloud project, add this SA in the new game's IAM page, and then give it the Firebase Admin role in both projects

B. Create a service account (SA) in the legacy game's Google Cloud project, add a second SA in the new game's IAM page, and then give the Organization Admin role to both SAs

C. Create a service account (SA) in the legacy game's Google Cloud project, give it the Firebase Admin role, and then migrate the new game to the legacy game's project.

D. Create a service account (SA) in the legacy game's Google Cloud project, give the SA the Organization Admin role and then give it the Firebase Admin role in both projects

Answer: B (LEAVE A REPLY)

Topic 10, Helicopter Racing League Case

Company Overview

Helicopter Racing League (HRL) is a global sports league for competitive helicopter racing. Each year HRL holds the world championship and several regional league competitions where teams compete to earn a spot in the world championship. HRL offers a paid service to stream the races all over the world with live telemetry and predictions throughout each race.

Solution concept

HRL wants to migrate their existing service to a new platform to expand their use of managed AI and ML services to facilitate race predictions. Additionally, as new fans engage with the sport, particularly in emerging regions, they want to move the serving of their content, both real-time and recorded, closer to their users.

Existing technical environment

HRL is a public cloud-first company; the core of their mission-critical applications runs on their current public cloud provider. Video recording and editing is performed at the race tracks, and the content is encoded and transcoded, where needed, in the cloud. Enterprise-grade connectivity and local compute is provided by truck-mounted mobile data centers. Their race prediction services are hosted exclusively on their existing public cloud provider. Their existing technical environment is as follows:

Existing content is stored in an object storage service on their existing public cloud provider.

Video encoding and transcoding is performed on VMs created for each job.

Race predictions are performed using TensorFlow running on VMs in the current public cloud provider.

Business Requirements

HRL's owners want to expand their predictive capabilities and reduce latency for their viewers in emerging markets. Their requirements are:

Support ability to expose the predictive models to partners.

Increase predictive capabilities during and before races:

- * Race results
- * Mechanical failures
- * Crowd sentiment

Increase telemetry and create additional insights.

Measure fan engagement with new predictions.

Enhance global availability and quality of the broadcasts.

Increase the number of concurrent viewers.

Minimize operational complexity.

Ensure compliance with regulations.

Create a merchandising revenue stream.

Technical Requirements

Maintain or increase prediction throughput and accuracy.

Reduce viewer latency.

Increase transcoding performance.

Create real-time analytics of viewer consumption patterns and engagement.

Create a data mart to enable processing of large volumes of race data.

Executive statement

Our CEO, S. Hawke, wants to bring high-adrenaline racing to fans all around the world. We listen to our fans, and they want enhanced video streams that include predictions of events within the race (e.g., overtaking). Our current platform allows us to predict race outcomes but lacks the facility to support real-time predictions during races and the capacity to process season-long results.

NEW QUESTION: 81

For this question, refer to the EHR Healthcare case study. EHR has single Dedicated Interconnect connection between their primary data center and Googles network. This connection satisfies EHR's network and security policies:

- * On-premises servers without public IP addresses need to connect to cloud resources without public IP addresses

- * Traffic flows from production network mgmt. servers to Compute Engine virtual machines should never traverse the public internet.

You need to upgrade the EHR connection to comply with their requirements. The new connection design must support business critical needs and meet the same network and security policy requirements. What should you do?

A. Add a new Dedicated Interconnect connection

B. Upgrade the bandwidth on the Dedicated Interconnect connection to 100 G

C. Add three new Cloud VPN connections

D. Add a new Carrier Peering connection

Answer: A ([LEAVE A REPLY](#))

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

NEW QUESTION: 82

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 ([LEAVE A REPLY](#))

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

NEW QUESTION: 83

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 ([LEAVE A REPLY](#))

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.

NEW QUESTION: 84

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: ([SHOW ANSWER](#)**)**

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: 85

For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the compute workloads for your company, Mountkirk Games.

Considering the Mountkirk Games business and technical requirements, what should you do?

A. Create a global load balancer with managed instance groups and autoscaling policies. Use preemptible Compute Engine instances.

B. Create a global load balancer with managed instance groups and autoscaling policies. Use non-preemptible Compute Engine instances.

C. Create network load balancers. Use preemptible Compute Engine instances.

D. Create network load balancers. Use non-preemptible Compute Engine instances.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 86

For this question, refer to the TerramEarth case study. TerramEarth has decided to store data files in Cloud Storage. You need to configure Cloud Storage lifecycle rule to store 1 year of data and minimize file storage cost.

Which two actions should you take?

A. Create a Cloud Storage lifecycle rule with Age: "30", Storage Class: "Standard", and Action: "Set to Coldline", and create a second GCS life-cycle rule with Age: "365", Storage Class: "Coldline", and Action: "Delete".

B. Create a Cloud Storage lifecycle rule with Age: "30", Storage Class: "Coldline", and Action: "Set to Nearline", and create a second GCS life-cycle rule with Age: "91", Storage Class: "Coldline", and Action: "Set to Nearline".

C. Create a Cloud Storage lifecycle rule with Age: "30", Storage Class: "Standard", and Action: "Set to Coldline", and create a second GCS life-cycle rule with Age: "365", Storage Class: "Nearline", and Action: "Delete".

D. Create a Cloud Storage lifecycle rule with Age: "90", Storage Class: "Standard", and Action: "Set to Nearline", and create a second GCS life-cycle rule with Age: "91", Storage Class: "Nearline", and Action: "Set to Coldline".

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 87

Mountkirk Games wants to limit the physical location of resources to their operating Google Cloud regions.

What should you do?

A. Configure an organizational policy which constrains where resources can be deployed.

B. Configure IAM conditions to limit what resources can be configured.

C. Configure the quotas for resources in the regions not being used to 0.

D. Configure a custom alert in Cloud Monitoring so you can disable resources as they are created in other

Answer: A ([LEAVE A REPLY](#))

regions.

NEW QUESTION: 88

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

A. Create a BigQuery table for the European data, and set the table retention period to 36 months. For Cloud Storage, use gsutil to enable lifecycle management using a DELETE action with an Age condition of 36 months.

B. Create a BigQuery table for the European data, and set the table retention period to 36 months. For Cloud Storage, use gsutil to create a SetStorageClass to NONE action when with an Age condition of 36 months.

C. Create a BigQuery time-partitioned table for the European data, and set the partition expiration period to 36 months. For Cloud Storage, use gsutil to enable lifecycle management using a DELETE action with an Age condition of 36 months.

D. Create a BigQuery time-partitioned table for the European data, and set the partition period to 36 months. For Cloud Storage, use gsutil to create a SetStorageClass to NONE action with an Age condition of 36 months.

Answer: ([SHOW ANSWER](#)**)**

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

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

NEW QUESTION: 89

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

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

Answer: ([SHOW ANSWER](#))

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

Google Compute Engine enables users to launch virtual machines (VMs) on demand. VMs can be launched from the standard images or custom images created by users.

Managed instance groups offer autoscaling capabilities that allow you to automatically add or remove instances from a managed instance group based on increases or decreases in load.

Autoscaling helps your applications gracefully handle increases in traffic and reduces cost when the need for resources is lower.

NEW QUESTION: 90

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

Mountkirk Games wants you to design their new testing strategy. How should the test coverage differ from their existing backends on the other platforms?

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

Answer: A ([LEAVE A REPLY](#))

From Scenario:

A few of their games were more popular than expected, and they had problems scaling their application servers, MySQL databases, and analytics tools.

Requirements for Game Analytics Platform include: Dynamically scale up or down based on game activity

Topic 3, TerramEarth

Company Overview

TerramEarth manufactures heavy equipment for the mining and agricultural industries: About 80% of their business is from mining and 20% from agriculture. They currently have over 500 dealers and service centers in 100 countries. Their mission is to build products that make their customers more productive.

Company Background


TerramEarth formed in 1946, when several small, family owned companies combined to retool after World War II. The company cares about their employees and customers and considers them to be extended members of their family.

TerramEarth is proud of their ability to innovate on their core products and find new markets as their customers' needs change. For the past 20 years trends in the industry have been largely toward increasing productivity by using larger vehicles with a human operator.

Solution Concept

There are 20 million TerramEarth vehicles in operation that collect 120 fields of data per second. Data is stored locally on the vehicle and can be accessed for analysis when a vehicle is serviced. The data is downloaded via a maintenance port. This same port can be used to adjust operational parameters, allowing the vehicles to be upgraded in the field with new computing modules. Approximately 200,000 vehicles are connected to a cellular network, allowing TerramEarth to collect data directly. At a rate of 120 fields of data per second, with 22 hours of operation per day. TerramEarth collects a total of about 9 TB/day from these connected vehicles.

Existing Technical Environment



TerramEarth's existing architecture is composed of Linux-based systems that reside in a data center. These systems gzip CSV files from the field and upload via FTP, transform and aggregate them, and place the data in their data warehouse. Because this process takes time, aggregated reports are based on data that is 3 weeks old.

With this data, TerramEarth has been able to preemptively stock replacement parts and reduce unplanned downtime of their vehicles by 60%. However, because the data is stale, some customers are without their vehicles for up to 4 weeks while they wait for replacement parts.

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 IP better position new products and services.
- * Have the ability to partner with different companies-especially with seed and fertilizer suppliers

CEO Statement We have been successful in capitalizing on the trend toward larger vehicles to increase the productivity of our customers. Technological change is occurring rapidly and

TerramEarth has taken advantage of connected devices technology to provide our customers with better services, such as our intelligent farming equipment. With this technology, we have been able to increase farmers' yields by 25%, by using past trends to adjust how our vehicles operate. These advances have led to the rapid growth of our agricultural product line, which we expect will generate 50% of our revenues by 2020.

CTO Statement

Our competitive advantage has always been in the manufacturing process with our ability to build better vehicles for lower cost than our competitors. However, new products with different approaches are constantly being developed, and I'm concerned that we lack the skills to undergo the next wave of transformations in our industry. Unfortunately, our CEO doesn't take technology obsolescence seriously and he considers the many new companies in our industry to be niche players. My goals are to build our skills while addressing immediate market needs through incremental innovations.

NEW QUESTION: 91

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

Answer: C (LEAVE A REPLY)

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

Valid Professional-Cloud-Architect Dumps shared by BraindumpsPass.com for Helping Passing Professional-Cloud-Architect Exam! BraindumpsPass.com now offer the **newest Professional-Cloud-Architect exam dumps**, the BraindumpsPass.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com Professional-Cloud-Architect dumps with Test Engine here: <https://www.braindumpsPASS.com/Google/Professional-Cloud-Architect-practice-exam-dumps.html> (378 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 92

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 services. The Stackdriver Logging agent must be installed on each instance to collect system logs.
- C.** Launch a custom syslogd compute instance and configure your GCP project and VMs to forward all logs to it.
- D.** Install the Stackdriver Logging agent on a single compute instance and let it collect all audit and access logs for your environment.

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 93

For this question, refer to the Helicopter Racing League (HRL) case study. Recently HRL started a new regional racing league in Cape Town, South Africa. In an effort to give customers in Cape Town a better user experience, HRL has partnered with the Content Delivery Network provider, Fastly. HRL needs to allow traffic coming from all of the Fastly IP address ranges into their Virtual Private Cloud network (VPC network). You are a member of the HRL security team and you need to configure the update that will allow only the Fastly IP address ranges through the External HTTP(S) load balancer. Which command should you use?

- A.** `gcloud compute firewall rules update hlr-policy \`
`--priority 1000 \`
`target tags-sourceiplist fastly \`
`--allow tcp:443`
- B.** `gcloud compute security policies rules update 1000 \`
`--security-policy hlr-policy \`
`--expression "evaluatePreconfiguredExpr('sourceiplist-fastly')"`
`--action "allow"`
- C.** `gcloud compute firewall rules update`
`sourceiplist-fastly \`
`priority 1000 \`
`allow tcp: 443`
- D.** `gcloud compute priority-policies rules update`
`1000 \`
`security policy from fastly`
`--src- ip-ranges"`

Answer: B ([LEAVE A REPLY](#))

-- action "allow"

Reference:

D18912E1457D5D1DDCBD40AB3BF70D5D

NEW QUESTION: 94

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

What should you do?

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

Answer: A ([LEAVE A REPLY](#))

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

NEW QUESTION: 95

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

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

Answer: ([SHOW ANSWER](#)**)**

<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: 96

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 (LEAVE A REPLY)

<https://cloud.google.com/serverless-options>

NEW QUESTION: 97

Your company operates nationally and plans to use GCP for multiple batch workloads, including some that are not time-critical. You also need to use GCP services that are HIPAA-certified and manage service costs.

How should you design to meet Google best practices?

- A. Provisioning preemptible VMs to reduce cost. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.
- B. Provisioning preemptible VMs to reduce cost. Disable and then discontinue use of all GCP and APIs that are not HIPAA-compliant.
- C. Provision standard VMs in the same region to reduce cost. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.
- D. Provision standard VMs to the same region to reduce cost. Disable and then discontinue use of all GCP services and APIs that are not HIPAA-compliant.

Answer: B (LEAVE A REPLY)

<https://cloud.google.com/security/compliance/hipaa/>

NEW QUESTION: 98

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

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

Answer: C (LEAVE A REPLY)

The framework illustrated in the preceding diagram has four phases:

* Assess. In this phase, you assess your source environment, assess the workloads that you want to migrate to Google Cloud, and assess which VMs support each workload.

- * Plan. In this phase, you create the basic infrastructure for Migrate for Compute Engine, such as provisioning the resource hierarchy and setting up network access.
- * Deploy. In this phase, you migrate the VMs from the source environment to Compute Engine.
- * Optimize. In this phase, you begin to take advantage of the cloud technologies and capabilities.

NEW QUESTION: 99

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: (SHOW ANSWER)

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: 100

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

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

Answer: C (LEAVE A REPLY)

<https://cloud.google.com/appengine/docs/standard/python3/connecting-vpc>

https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databases#on_premises

NEW QUESTION: 101

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

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

Answer: ([SHOW ANSWER](#))

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

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

NEW QUESTION: 102

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 103

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

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

Answer: A,C ([LEAVE A REPLY](#))

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

NEW QUESTION: 104

You want to enable your running Google Container Engine cluster to scale as demand for your application changes.

What should you do?

A. Add additional nodes to your Container Engine cluster using the following command:

```
gcloud container clusters resize CLUSTER_NAME --size 10
```

B. Add a tag to the instances in the cluster with the following command:

```
gcloud compute instances add-tags INSTANCE --tags enable --autoscaling max-nodes-10
```

C. Update the existing Container Engine cluster with the following command:

```
gcloud alpha container clusters update mycluster --enable-autoscaling --min-nodes=1 --max-nodes=10
```

D. Create a new Container Engine cluster with the following command:

```
gcloud alpha container clusters create mycluster --enable-autocaling --min-nodes=1 --max-nodes=10
```

Answer: B ([LEAVE A REPLY](#))

and redeploy your application.

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler> Cluster autoscaling
--enable-autoscaling

Enables autoscaling for a node pool.

Enables autoscaling in the node pool specified by --node-pool or the default node pool if --node-pool is not provided.

Where:

--max-nodes=MAX_NODES

Maximum number of nodes in the node pool.

Maximum number of nodes to which the node pool specified by --node-pool (or default node pool if unspecified) can scale.

Incorrect Answers:

C, D: Warning: Do not use Alpha Clusters or alpha features for production workloads.

Note: You can experiment with Kubernetes alpha features by creating an alpha cluster. Alpha clusters are short-lived clusters that run stable Kubernetes releases with all Kubernetes APIs and features enabled. Alpha clusters are designed for advanced users and early adopters to experiment with workloads that take advantage of new features before those features are production-ready. You can use Alpha clusters just like normal Kubernetes Engine clusters.

NEW QUESTION: 105

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.pr
    At
    Sun.security.util.manifestEntryVerifier.ver
    At java.net.URLClassLoader.defineClass
    At sun.reflect.GeneratedMethodAccessor
    At
    Sun.reflect.DelegatingMethodAccessorImpl.
    At java.lang.reflect.Method.invoke
```

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

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 106

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. In the GCP Console, navigate to Stackdriver Logging. Consult logs for Kubernetes Engine and Cloud SQL.
- B. Use `gcloud sql instances restart`.
- C. In the GCP Console, navigate to Cloud SQL. Restore the latest backup. Use `kubect1` to restart all pods.
- D. Validate that the Service Account used by the Cloud SQL proxy container still has the Cloud Build Editor role.

Answer: A ([LEAVE A REPLY](#))

Valid Professional-Cloud-Architect Dumps shared by BraindumpsPass.com for Helping Passing Professional-Cloud-Architect Exam! BraindumpsPass.com now offer the **newest**

Professional-Cloud-Architect exam dumps, the BraindumpsPass.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com Professional-Cloud-Architect dumps with Test Engine here: <https://www.braindumpsPass.com/Google/Professional-Cloud-Architect-practice-exam-dumps.html> (378 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 107

For this question, refer to the JencoMart case study.

JencoMart has decided to migrate user profile storage to Google Cloud Datastore and the application servers to Google Compute Engine (GCE). During the migration, the existing infrastructure will need access to Datastore to upload the data. What service account key-management strategy should you recommend?

- A.** Provision service account keys for the on-premises infrastructure and for the GCE virtual machines (VMs).
- B.** Authenticate the on-premises infrastructure with a user account and provision service account keys for the VMs.
- C.** Provision service account keys for the on-premises infrastructure and use Google Cloud Platform (GCP) managed keys for the VMs
- D.** Deploy a custom authentication service on GCE/Google Container Engine (GKE) for the on-premises infrastructure and use GCP managed keys for the VMs.

Answer: A ([LEAVE A REPLY](#))

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

Migrating data to Google Cloud Platform

Let's say that you have some data processing that happens on another cloud provider and you want to transfer the processed data to Google Cloud Platform. You can use a service account from the virtual machines on the external cloud to push the data to Google Cloud Platform. To do this, you must create and download a service account key when you create the service account and then use that key from the external process to call the Cloud Platform APIs.

Reference:

https://cloud.google.com/iam/docs/understanding-service-accounts#migrating_data_to_google_cloud_platform

NEW QUESTION: 108

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.** Enable autoscaling on the GKE cluster and enable horizontal pod autoscaling on your application deployments. Send curl requests to your application, and validate if the auto scaling works.

- B.** Use a load testing tool to simulate the expected number of concurrent users and total requests to your application, and inspect the results.
- C.** 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.
- D.** Use Cloud Debugger in the development environment to understand the latency between the different microservices.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 109

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

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

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 110

For this question, refer to the Helicopter Racing League (HRL) case study. A recent finance audit of cloud infrastructure noted an exceptionally high number of Compute Engine instances are allocated to do video encoding and transcoding. You suspect that these Virtual Machines are zombie machines that were not deleted after their workloads completed. You need to quickly get a list of which VM instances are idle. What should you do?

- A.** From the Google Cloud Console, identify which Compute Engine instances in the managed instance groups are no longer responding to health check probes.
- B.** Use the `gcloud recomender` command to list the idle virtual machine instances.
- C.** Log into each Compute Engine instance and collect disk, CPU, memory, and network usage statistics for analysis.
- D.** Use the `gcloud compute instances list` to list the virtual machine instances that have the `idle: true` label set.

Answer: (SHOW ANSWER)

NEW QUESTION: 111

Your web application uses Google Kubernetes Engine to manage several workloads. One workload requires a consistent set of hostnames even after pod scaling and relaunches.

Which feature of Kubernetes should you use to accomplish this?

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

Answer: A ([LEAVE A REPLY](#))

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

NEW QUESTION: 112

For this question, refer to the TerramEarth case study

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

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

Answer: A ([LEAVE A REPLY](#))

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

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

Reference:

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

NEW QUESTION: 113

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

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

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 114

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

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

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

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

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 115

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

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 116

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

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

- A. Local SSD for customer session state data. Lifecycle-managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.
- B. Memcache backed by Cloud Datastore for the customer session state data. Lifecycle-managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.
- C. Memcache backed by Cloud SQL for customer session state data. Assorted local SSD-backed instances for VM boot/data volumes. Cloud Storage for log archives and thumbnails.

D. Memcache backed by Persistent Disk SSD storage for customer session state data. Assorted local SSD-backed instances for VM boot/data volumes. Cloud Storage for log archives and thumbnails.

Answer: D ([LEAVE A REPLY](#))

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

NEW QUESTION: 117

For this question, refer to the Dress4Win case study.

At Dress4Win, an operations engineer wants to create a low-cost solution to remotely archive copies of database backup files. The database files are compressed tar files stored in their current data center. How should he proceed?

- A.** Create a cron script using gsutil to copy the files to a Coldline Storage bucket.
- B.** Create a cron script using gsutil to copy the files to a Regional Storage bucket.
- C.** Create a Cloud Storage Transfer Service Job to copy the files to a Coldline Storage bucket.
- D.** Create a Cloud Storage Transfer Service job to copy the files to a Regional Storage bucket.

Answer: A ([LEAVE A REPLY](#))

Follow these rules of thumb when deciding whether to use gsutil or Storage Transfer Service:

- * When transferring data from an on-premises location, use gsutil.
- * When transferring data from another cloud storage provider, use Storage Transfer Service.
- * Otherwise, evaluate both tools with respect to your specific scenario.

Use this guidance as a starting point. The specific details of your transfer scenario will also help you determine which tool is more appropriate

<https://cloud.google.com/storage-transfer/docs/overview>

NEW QUESTION: 118

You have found an error in your App Engine application caused by missing Cloud Datastore indexes. You have created a YAML file with the required indexes and want to deploy these new indexes to Cloud Datastore.

What should you do?

- A.** Create an HTTP request to the built-in python module to send the index configuration file to your application
- B.** Point gcloud datastore create-indexes to your configuration file
- C.** In the GCP Console, use Datastore Admin to delete the current indexes and upload the new configuration file
- D.** Upload the configuration file to the App Engine's default Cloud Storage bucket, and have App Engine detect the new indexes

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 119

Your company has an application running as a Deployment in a Google Kubernetes Engine (GKE) cluster. When releasing new versions of the application via a rolling deployment, the team

has been causing outages The root cause of the outages is misconfigurations with parameters that are only used in production You want to put preventive measures for this in the platform to prevent outages What should you do?

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

Answer: A ([LEAVE A REPLY](#))

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

<https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps>

<https://cloud.google.com/blog/products/containers-kubernetes/kubernetes-best-practices-setting-up-health-checks-with-readiness-and-liveness-probes>

NEW QUESTION: 120

For this question, refer to the Helicopter Racing League (HRL) case study. The HRL development team releases a new version of their predictive capability application every Tuesday evening at 3 a.m. UTC to a repository. The security team at HRL has developed an in-house penetration test Cloud Function called Airwolf.

The security team wants to run Airwolf against the predictive capability application as soon as it is released every Tuesday. You need to set up Airwolf to run at the recurring weekly cadence. What should you do?

- A. Set up Cloud Tasks and a Cloud Storage bucket that triggers a Cloud Function.
- B. Set up a Cloud Logging sink and a Cloud Storage bucket that triggers a Cloud Function.
- C. Configure the deployment job to notify a Pub/Sub queue that triggers a Cloud Function.
- D. Set up Identity and Access Management (IAM) and Confidential Computing to trigger a Cloud Function.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 121

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: ([SHOW ANSWER](#))

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). Reference:

<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>

Valid Professional-Cloud-Architect Dumps shared by BraindumpsPass.com for Helping Passing Professional-Cloud-Architect Exam! BraindumpsPass.com now offer the **newest Professional-Cloud-Architect exam dumps**, the BraindumpsPass.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com Professional-Cloud-Architect dumps with Test Engine here: <https://www.braindumpsPass.com/Google/Professional-Cloud-Architect-practice-exam-dumps.html> (378 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 122

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

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 123

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

Answer: A ([LEAVE A REPLY](#))

an incident.

Explanation:

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

NEW QUESTION: 124

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

Answer: C,E ([LEAVE A REPLY](#))

Valid Professional-Cloud-Architect Dumps shared by BraindumpsPass.com for Helping Passing Professional-Cloud-Architect Exam! BraindumpsPass.com now offer the **newest Professional-Cloud-Architect exam dumps**, the BraindumpsPass.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com Professional-Cloud-Architect dumps with Test Engine here: <https://www.braindumpsPass.com/Google/Professional-Cloud-Architect-practice-exam-dumps.html> (378 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)