Google.Professional-Cloud-Architect.v2024-01-19.q131

Exam Code:	Professional-Cloud-Architect
Exam Name:	Google Certified Professional - Cloud Architect (GCP)
Certification Provider:	Google
Free Question Number:	131
Version:	v2024-01-19
# of views:	1104
# of Questions views:	1310
https://www.freepdfdumps.com/Google.Professional-Cloud-Architect.v2024-01-19.q131.html	

NEW QUESTION: 1
You are working at a financial institution that stores mortgage loan approval documents on Cloud Storage.
Any change to these approval documents must be uploaded as a separate approval file, so you want to ensure that these documents cannot be deleted or overwritten for the next 5 years. What should you do?
A. Create a retention policy on the bucket for the duration of 5 years. Create a lock on the retention policy.
B. Create the bucket with uniform bucket-level access, and grant a service account the role of Object Writer. Use the service account to upload new files.
C. Use a customer-managed key for the encryption of the bucket. Rotate the key after 5 years.
D. Create the bucket with fine-grained access control, and grant a service account the role of Object Writer.Use the service account to upload new files.
Answer: A (LEAVE A REPLY)
Reference: https://cloud.google.com/storage/docs/using-bucket-lock

NEW QUESTION: 2
For this question, refer to the TerramEarth case study
Your development team has created a structured API to retrieve vehicle data. They want to allow third parties to develop tools for dealerships that use this vehicle event data. You want to support delegated authorization against this data. What should you do?
A. Build or leverage an OAuth-compatible access control system.
B. Build SAML 2.0 SSO compatibility into your authentication system.
C. Restrict data access based on the source IP address of the partner systems.
D. Create secondary credentials for each dealer that can be given to the trusted third party.
Answer: A (LEAVE A REPLY)
Explanation
https://cloud.google.com/appengine/docs/flexible/go/authorizing-apps

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

NEW QUESTION: 3
Your web application must comply with the requirements of the European Union's General Data Protection Regulation (GDPR). You are responsible for the technical architecture of your web application. What should you do?
A. Ensure that your web application only uses native features and services of Google Cloud Platform, because Google already has various certifications and provides "pass-on" compliance when you use native features.
B. Enable the relevant GDPR compliance setting within the GCPConsole for each of the services in use within your application.
C. Ensure that Cloud Security Scanner is part of your test planning strategy in order to pick up any compliance gaps.
D. Define a design for the security of data in your web application that meets GDPR requirements.
Answer: (SHOW ANSWER)
Explanation
https://cloud.google.com/security/gdpr/?tab=tab4
Reference: https://www.mobiloud.com/blog/gdpr-compliant-mobile-app/

NEW QUESTION: 4
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. Send the replicated transaction to Google Cloud Pub/Sub.
C. Add additional VPN connections and load balance them.
D. Restore their database daily using Google Cloud SQL.
E. Configure their replication to use UDP.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 5

A development team at your company has created a dockerized HTTPS web application. You need to deploy the application on Google Kubernetes Engine (GKE) and make sure that the application scales automatically.
How should you deploy to GKE?
A. Use the Horizontal Pod Autoscaler and enable cluster 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)
Explanation
https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer
https://cloud.google.com/kubernetes-engine/docs/concepts/network-overview#ext-lb

NEW QUESTION: 6
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 Vision AI.
C. Use Google Cloud's operations suite.
D. Use Jupyter Notebooks.
Answer: A (LEAVE A REPLY)
Reference: https://cloud.google.com/ai-platform/prediction/docs/ai-explanations/preparing-metadata
Topic 9, EHR Health Care
Company overview
EHR Healthcare is a leading provider of electronic health record software to the medical industry. EHR Healthcare provides their software as a service to multi-national medical offices, hospitals, and insurance providers.
Solution concept
Due to rapid changes in the healthcare and insurance industry, EHR Healthcare's business has been growing exponentially year over year. They need to be able to scale their environment, adapt their disaster recovery plan, and roll out new continuous deployment capabilities to update their software at a fast pace. Google Cloud has been chosen to replace their current colocation facilities.
Existing technical environment
EHR's software is currently hosted in multiple colocation facilities. The lease on one of the data centers is about to expire.

Customer-facing applications are web-based, and many have recently been containerized to run on a group of Kubernetes clusters. Data is stored in a mixture of relational and NoSQL databases (MySQL, MS SQL Server, Redis, and MongoDB).
EHR is hosting several legacy file- and API-based integrations with insurance providers on-premises. These systems are scheduled to be replaced over the next several years. There is no plan to upgrade or move these systems at the current time.
Users are managed via Microsoft Active Directory. Monitoring is currently being done via various open source tools. Alerts are sent via email and are often ignored.
Business requirements
* On-board new insurance providers as quickly as possible.
* Provide a minimum 99.9% availability for all customer-facing systems.
* Provide centralized visibility and proactive action on system performance and usage.
* Increase ability to provide insights into healthcare trends.
* Reduce latency to all customers.
* Maintain regulatory compliance.
* Decrease infrastructure administration costs.
* Make predictions and generate reports on industry trends based on provider data.
Technical requirements
* Maintain legacy interfaces to insurance providers with connectivity to both on-premises systems and cloud providers.
* Provide a consistent way to manage customer-facing applications that are container-based.
* Provide a secure and high-performance connection between on-premises systems and Google Cloud.
* Provide consistent logging, log retention, monitoring, and alerting capabilities.
* Maintain and manage multiple container-based environments.
* Dynamically scale and provision new environments.
* Create interfaces to ingest and process data from new providers.
Executive statement
Our on-premises strategy has worked for years but has required a major investment of time and money in training our team on distinctly different systems, managing similar but separate environments, and responding to outages. Many of these outages have been a result of misconfigured systems, inadequate capacity to manage spikes in traffic, and inconsistent monitoring practices. We want to use Google Cloud to leverage a scalable, resilient platform that can span multiple environments seamlessly and provide a consistent and stable user experience that positions us for future growth.

NEW QUESTION: 7
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: A (LEAVE A REPLY)
Explanation
https://cloud.google.com/appengine/docs/standard/python3/connecting-vpc
https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databases#on_premises

NEW QUESTION: 8
Your company has a Google Cloud project that uses BigQuery for data warehousing They have a VPN tunnel between the on-premises environment and Google Cloud that is configured with Cloud VPN. The security team wants to avoid data exfiltration by malicious insiders, compromised code, and accidental oversharing.
What should they do?
A. Configure Private Google Access for on-premises only.
B. Perform the following tasks:
1) Create a service account.
2) Give the BigQuery JobUser role and Storage Reader role to the service account.
3) Remove all other IAM access from the project.
C. Configure VPC Service Controls and configure Private Google Access.
D. Configure Private Google Access.
Answer: C (LEAVE A REPLY)
Explanation
https://cloud.google.com/vpc-service-controls/docs/overview
VPC Service Controls improves your ability to mitigate the risk of data exfiltration from Google Cloud services such as Cloud Storage and BigQuery.

NEW QUESTION: 9
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: D (LEAVE A REPLY)
Explanation
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). References:
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-

NEW QUESTION: 10
You are managing several internal applications that are deployed on Compute Engine. Business users inform you that an application has become very slow over the past few days. You want to find the underlying cause in order to solve the problem. What should you do first?
A. Inspect the logs and metrics from the instances in Cloud Logging and Cloud Monitoring.
B. Restore a backup of the application database from a time before the application became slow.
C. Deploy the applications on a managed instance group with autoscaling enabled. Add a load balancer in front of the managed instance group, and have the users connect to the IP of the load balancer.
D. Change the Compute Engine Instances behind the application to a machine type with more CPU and memory.
Answer: A (LEAVE A REPLY)
Explanation
When an application becomes slow, the first step you should take is to gather information about the underlying cause of the problem. One way to do this is by inspecting the logs and metrics from the instances where the application is deployed. Google Cloud Platform (GCP) provides tools such as Cloud Logging and Cloud Monitoring that can help you to collect and analyze this information. By reviewing the logs and metrics from the instances, you may be able to identify issues such as resource shortages (e.g. CPU, memory, or disk), network problems, or application errors that are causing the performance issues. Once you have identified the underlying cause of the problem, you can take steps to resolve it.

NEW QUESTION: 11
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)
Explanation

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

NEW QUESTION: 12
Your company is designing its application landscape on Compute Engine. Whenever a zonal outage occurs, the application should be restored in another zone as quickly as possible with the latest application data. You need to design the solution to meet this requirement. What should you do?
A. Create a snapshot schedule for the disk containing the application 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: B (LEAVE A REPLY)
Explanation
Regional persistent disk is a storage option that provides synchronous replication of data between two zones in a region. Regional persistent disks can be a good building block to use when you implement HA services in Compute Engine. https://cloud.google.com/compute/docs/disks/high-availability-regional-persistent-disk

NEW QUESTION: 13

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

NEW QUESTION: 14
You created a pipeline that can deploy your source code changes to your infrastructure in instance groups for self healing.
One of the changes negatively affects your key performance indicator.
You are not sure how to fix it and investigation could take up to a week.
What should you do
A. Log into the servers with the bad code change, and swap in the previous code
B. Log in to a server, and iterate a fix locally
C. Change the instance group template to the previous one, and delete all instances.
D. Revert the source code change and rerun the deployment pipeline
Answer: D (LEAVE A REPLY)

NEW QUESTION: 15
For this question, refer to the Mountkirk Games case study
Mountkirk Games needs to create a repeatable and configurable mechanism for deploying isolated application environments. Developers and testers can access each other's environments and resources, but they cannot access staging or production resources. The staging environment needs access to some services from production.
What should you do to isolate development environments from staging and production?
A. Create a project for development and test and another for staging and production.
B. Create one project for development, a second for staging and a third for production.

C. Create a network for development and test and another for staging and production.
D. Create one subnetwork for development and another for staging and production.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 16
For this question, refer to the TerramEarth case study.
TerramEarth's 20 million vehicles are scattered around the world. Based on the vehicle's location its telemetry data is stored in a Google Cloud Storage (GCS) regional bucket (US. Europe, or Asia). The CTO has asked you to run a report on the raw telemetry data to determine why vehicles are breaking down after 100 K miles.
You want to run this job on all the data. What is the most cost-effective way to run this job?
A. Move all the data into 1 zone, then launch a Cloud Dataproc cluster to run the job.
B. Move all the data into 1 region, then launch a Google Cloud Dataproc cluster to run the job.
C. Launch a cluster in each region to preprocess and compress the raw data, then move the data into a multi region bucket and use a Dataproc cluster to finish the job.
D. Launch a cluster in each region to preprocess and compress the raw data, then move the data into a region bucket and use a Cloud Dataproc cluster to finish the jo
Answer: D (LEAVE A REPLY)
Explanation
Storageguarantees 2 replicates which are geo diverse (100 miles apart) which can get better remote latency and availability.
More importantly, is that multiregional heavily leverages Edge caching and CDNs to provide the content to the end users.
All this redundancy and caching means that Multiregional comes with overhead to sync and ensure consistency between geo-diverse areas. As such, it's much better for write-once-read-many scenarios.
This means frequently accessed (e.g. "hot" objects) around the world, such as website content, streaming videos, gaming or mobile applications.
References:
https://medium.com/google-cloud/google-cloud-storage-what-bucket-class-for-the-best-performance-5c847ac8f9
Topic 2, Mountkirk Games Case Study 1
Company Overview
Mountkirk Games makes online, session-based. multiplayer games for the most popular mobile platforms.
Company Background
Mountkirk Games builds all of their games with some server-side integration and has historically used cloud providers to lease physical servers. A few of their games were more popular than expected, and they had problems scaling their application servers, MySQL databases, and analytics tools.
Mountkirk's 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.
Technical Requirements
Requirements for Game Backend Platform
1. Dynamically scale up or down based on game activity.
2. Connect to a managed NoSQL database service.
3. Run customized Linx distro.
Requirements for Game Analytics Platform
1. Dynamically scale up or down based on game activity.
2. Process incoming data on the fly directly from the game servers.
3. Process data that arrives late because of slow mobile networks.
4. Allow SQL queries to access at least 10 TB of historical data.
5. Process files that are regularly uploaded by users' mobile devices.
6. Use only fully managed services
CEO Statement
Our last successful game did not scale well with our previous cloud provider, resuming 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 gams to target users.
CTO Statement
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.
CFO Statement
We are not capturing enough user demographic data usage metrics, and other KPIs. As a result, we do not engage the right users. We are not confident that our marketing is targeting the right users, and we are not selling enough premium Blast-Ups inside the games, which dramatically impacts our revenue.

NEW QUESTION: 17

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

A. Delete records older than 30 days.

B. Double the number of nodes you currently have.

C. Advise your clients to use HBase APIs instead of NodeJS APIs.

D. Review your RowKey strategy and ensure that keys are evenly spread across the alphabet.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 18

Your company has just acquired another company, and you have been asked to integrate their existing Google Cloud environment into your company's data center. Upon investigation, you discover that some of the RFC

1918 IP ranges being used in the new company's Virtual Private Cloud (VPC) overlap with your data center IP space. What should you do to enable connectivity and make sure that there are no routing conflicts when connectivity is established?

A. Create a Cloud VPN connection from the new VPC to the data center, create a Cloud Router, and apply new IP addresses so there is no overlapping IP space.

B. Create a Cloud VPN connection from the new VPC to the data center, and create a Cloud NAT instance to perform NAT on the overlapping IP space.

C. Create a Cloud VPN connection from the new VPC to the data center, create a Cloud Router, and apply a custom route advertisement to block the overlapping IP space.

D. Create a Cloud VPN connection from the new VPC to the data center, and apply a firewall rule that blocks the overlapping IP space.

Answer: (SHOW ANSWER)

Explanation

To connect two networks together we need (1) either VPN or interconnect and (2) peering. When there is peering, you cannot have conflicting IP addresses. You can use either Cloud VPN or Cloud Interconnect to securely connect your on-premises network to your VPC network.

(https://cloud.google.com/vpc/docs/vpc-peering#transit-network) At the time of peering, Google Cloud checks to see if there are any subnet IP ranges that overlap subnet IP ranges in the other network. If there is any overlap, peering is not established. (https://cloud.google.com/vpc/docs/vpc-peering#considerations) NAT is used to translate private to public IP and vice versa, however because we are connecting 2 networks together, they become private IPs. So it is not applicable.

NEW QUESTION: 19

You are managing several projects on Google Cloud and need to interact on a daily basis with BigQuery, Bigtable and Kubernetes Engine using the gcloud CLI tool You are travelling a lot and work on different

workstations during the week You want to avoid having to manage the gcloud CLI manually What should you do?

A. Use a package manager to install gcloud on your workstations instead of installing it manually

B. Create a Compute Engine instance and install gcloud on the instance Connect to this instance via SSH to always use the same gcloud installation when interacting with Google Cloud

C. Install gcloud on all of your workstations Run the command gcloud components auto-update on each workstation

D. Use Google Cloud Shell in the Google Cloud Console to interact with Google Cloud

Answer: D (LEAVE A REPLY)

Explanation

This option allows you to use the gcloud CLI tool without having to install or manage it manually on different workstations. Google Cloud Shell is a browser-based command-line tool that provides you with a temporary Compute Engine virtual machine instance preloaded with the Cloud SDK, including the gcloud CLI tool. You can access Google Cloud Shell from any web browser and use it to interact with BigQuery, Bigtable and Kubernetes Engine using the gcloud CLI tool. The other options are not optimal for this scenario, because they either require installing and updating the gcloud CLI tool on multiple workstations (A, C), or creating and maintaining a Compute Engine instance for the sole purpose of using the gcloud CLI tool (B). References:

https://cloud.google.com/shell/docs/overview

https://cloud.google.com/sdk/gcloud/

NEW QUESTION: 20

You need to develop procedures to test a disaster plan for a mission-critical application. You want to use Google-recommended practices and native capabilities within GCP.

What should you do?

A. Use Deployment Manager to automate service provisioning. Use Activity Logs to monitor and debug your tests.

B. Use Deployment Manager to automate provisioning. Use Stackdriver to monitor and debug your tests.

C. Use gcloud scripts to automate service provisioning. Use Activity Logs monitor and debug your tests.

D. Use automated scripts to automate service provisioning. Use Activity Logs monitor and debug your tests.

Answer: (SHOW ANSWER)

Explanation

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

NEW QUESTION: 21

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)
Explanation
https://cloud.google.com/datastore/docs/concepts/overview
Common workloads for Google Cloud Datastore:
* User profiles
* Product catalogs
* Game state
References: https://cloud.google.com/storage-options/
https://cloud.google.com/datastore/docs/concepts/overview

NEW QUESTION: 22
You are migrating a Linux-based application from your private data center to Google Cloud. The TerramEarth security team sent you several recent Linux vulnerabilities published by Common Vulnerabilities and Exposures (CVE). You need assistance in understanding how these vulnerabilities could impact your migration. What should you do?
A. Open a support case regarding the CVE and chat with the support engineer.
B. Read the CVEs from the Google Cloud Status Dashboard to understand the impact.
C. Read the CVEs from the Google Cloud Platform Security Bulletins to understand the impact
D. Post a question regarding the CVE in Stack Overflow to get an explanation
E. Post a question regarding the CVE in a Google Cloud discussion group to get an explanation
Answer: A,C (LEAVE A REPLY)
Explanation
https://cloud.google.com/support/bulletins

NEW QUESTION: 23
You are working with a data warehousing team that performs data analysis. The team needs to process data from external partners, but the data contains personally identifiable information (PII). You need to process and store the data without storing any of the Pll data. What should you do?
A. Create a Dataflow pipeline to retrieve the data from the external sources. As part of the pipeline use the Cloud Data Loss Prevention (Cloud DLP) API to remove any Pll data Store the result in BigQuery
B. Create a Dataflow pipeline to retrieve the data from the external sources. As part of the pipeline store all non-PII data in BigQuery and store all Pll data in a Cloud Storage bucket that has a retention policy set.
C. Ask the external partners to upload an data on Cloud Storage Configure Bucket Lock for the bucket Create a Dataflow pipeline to read the data from the bucket As part of the pipeline, use the Cloud Data Loss Prevention (Cloud DIP) API to remove any Pll data Store the result in BigQuery
D. Ask the external partners to import ail data in your BigQuery dataset Create a dataflow pipeline to copy the data into a new table As part of the Dataflow bucket skip all data in columns that have Pll data
Answer: A (LEAVE A REPLY)

Explanation
Create a Dataflow pipeline to retrieve the data from the external sources, he did not specify the way he is going to create it, it might be a pub/sub or external table or whatever.

NEW QUESTION: 24
Your organization has decided to restrict the use of external IP addresses on instances to only approved instances. You want to enforce this requirement across all of your Virtual Private Clouds (VPCs). What should you do?
A. Remove the default route on all VPCs. Move all approved instances into a new subnet that has a default route to an internet gateway.
B. Create a new VPC in custom mode. Create a new subnet for the approved instances, and set a default route to the internet gateway on this new subnet.
C. Implement a Cloud NAT solution to remove the need for external IP addresses entirely.
D. Set an Organization Policy with a constraint on constraints/compute.vmExternalIpAccess. List the approved instances in the allowedValues list.
Answer: (SHOW ANSWER)
Reference: https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#disableexternalip you might want to restrict external IP address so that only specific VM instances can use them. This option can help to prevent data exfiltration or maintain network isolation. Using an Organization Policy, you can restrict external IP addresses to specific VM instances with constraints to control use of external IP addresses for your VM instances within an organization or a project.

NEW QUESTION: 25
You need to deploy a stateful workload on Google Cloud. The workload can scale horizontally, but each instance needs to read and write to the same POSIX filesystem. At high load, the stateful workload needs to support up to 100 MB/s of writes. What should you do?
A. Use a persistent disk for each instance.
B. Use a regional persistent disk for each instance.
C. Create a Cloud Filestore instance and mount it in each instance.
D. Create a Cloud Storage bucket and mount it in each instance using gcsfuse.
Answer: C (LEAVE A REPLY)
Explanation
https://cloud.google.com/storage/docs/gcs-fuse#notes
Cloud Filestore: Cloud Filestore is a scalable and highly available shared file service fully managed by Google. Cloud Filestore provides persistent storage ideal for shared workloads. It is best suited for enterprise applications requiring persistent, durable, shared storage which is accessed by NFS or requires a POSIX compliant file system.
Reference: https://cloud.google.com/storage/docs/gcs-fuse

NEW QUESTION: 26
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)
Explanation
https://cloud.google.com/storage/docs/lifecycle

NEW QUESTION: 27
You need to set up Microsoft SQL Server on GCP. Management requires that there's no downtime in case of a data center outage in any of the zones within a GCP region. What should you do?
A. Configure a Cloud SQL instance with high availability enabled.
B. Configure a Cloud Spanner instance with a regional instance configuration.
C. Set up SQL Server on Compute Engine, using Always On Availability Groups using Windows Failover Clustering. Place nodes in different subnets.
D. Set up SQL Server Always On Availability Groups using Windows Failover Clustering. Place nodes in different zones.
Answer: D (LEAVE A REPLY)
Explanation
https://cloud.google.com/sql/docs/sqlserver/configure-ha

NEW QUESTION: 28
You are analyzing and defining business processes to support your startup's trial usage of GCP, and you don't yet know what consumer demand for your product will be. Your manager requires you to minimize GCP service costs and adhere to Google best practices. What should you do?
A. Utilize free tier and sustained use discounts. Provision a staff position for service cost management.
B. Utilize free tier and sustained use discounts. Provide training to the team about service cost management.
C. Utilize free tier and committed use discounts. Provision a staff position for service cost management.
D. Utilize free tier and committed use discounts. Provide training to the team about service cost management.
Answer: (SHOW ANSWER)
Explanation
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#billing_and_management

NEW QUESTION: 29
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)
Explanation
https://cloud.google.com/interconnect/docs/how-to/direct-peering

NEW QUESTION: 30
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. Install the Stackdriver agent on all of the legacy web servers.
B. In the Cloud Platform Console download the list of the uptime servers' IP addresses and create an inbound firewall rule
C. Configure their load balancer to pass through the User-Agent HTTP header when the value matches GoogleStackdriverMonitoring-UptimeChecks (https://cloud.google.com/monitoring)
D. 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)
Answer: B (LEAVE A REPLY)

NEW QUESTION: 31
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)

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

NEW QUESTION: 32
You are developing your microservices application on Google Kubernetes Engine. During testing, you want to validate the behavior of your application in case a specific microservice should suddenly crash. What should you do?
A. Add a taint to one of the nodes of the Kubernetes cluster. For the specific microservice, configure a pod anti-affinity label that has the name of the tainted node as a value.
B. Use Istio's fault injection on the particular microservice whose faulty behavior you want to simulate.
C. Destroy one of the nodes of the Kubernetes cluster to observe the behavior.
D. Configure Istio's traffic management features to steer the traffic away from a crashing microservice.
Answer: B (LEAVE A REPLY)
Explanation
Microservice runs on all nodes. The Micro service runs on Pod, Pod runs on Nodes. Nodes is nothing but Virtual machines. Once deployed the application microservices will get deployed across all Nodes. Destroying one node may not mimic the behaviour of microservice crashing as it may be running in other nodes.
link: https://istio.io/latest/docs/tasks/traffic-management/fault-injection/

NEW QUESTION: 33
Your operations team has asked you to help diagnose a performance issue in a production application that runs on Compute Engine. The application is dropping requests that reach it when under heavy load. The process list for affected instances shows a single application process that is consuming all available CPU, and autoscaling has reached the upper limit of instances. There is no abnormal load on any other related systems, including the database. You want to allow production traffic to be served again as quickly as possible. Which action should you recommend?
A. Change the autoscaling metric to agent.googleapis.com/memory/percent_used.
B. Restart the affected instances on a staggered schedule.
C. SSH to each instance and restart the application process.
D. Increase the maximum number of instances in the autoscaling group.
Answer: (SHOW ANSWER)

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

NEW QUESTION: 34
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: "Nearline", and Action: "Delete".
B. 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".
C. 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".
D. 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".
Answer: B (LEAVE A REPLY)

NEW QUESTION: 35
Your development teams release new versions of games running on Google Kubernetes Engine (GKE) daily.
You want to create service level indicators (SLIs) to evaluate the quality of the new versions from the user's perspective. What should you do?
A. Create Request Latency and Error Rate as service level indicators.
B. Create CPU Utilization and Request Latency as service level indicators.
C. Create GKE CPU Utilization and Memory Utilization as service level indicators.
D. Create Server Uptime and Error Rate as service level indicators.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 36
Your development team has created a mobile game app. You want to test the new mobile app on Android and iOS devices with a variety of configurations. You need to ensure that testing is efficient and cost-effective.
What
should you do?
A. Create Android and iOS VMs on Google Cloud, install the mobile app on the VMs, and test the mobile app.

B. Create Android and iOS containers on Google Kubernetes Engine (GKE), install the mobile app on the containers, and test the mobile app.
C. Upload your mobile app with different configurations to Firebase Hosting and test each configuration.
D. Upload your mobile app to the Firebase Test Lab, and test the mobile app on Android and iOS devices.
Answer: (SHOW ANSWER)

NEW QUESTION: 37
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)
Explanation
https://cloud.google.com/appengine/docs/standard/php/memcache/using

NEW QUESTION: 38
Your company has a networking team and a development team. The development team runs applications on Compute Engine instances that contain sensitive data. The development team requires administrative permissions for Compute Engine. Your company requires all network resources to be managed by the networking team. The development team does not want the networking team to have access to the sensitive data on the instances. What should you do?
A. 1. Create a project with a standalone VPC and assign the Network Admin role to the networking team.
2.Create a second project with a standalone VPC and assign the Compute Admin role to the development team.
3.Use Cloud VPN to join the two VPCs.
B. 1. Create a project with a standalone Virtual Private Cloud (VPC), assign the Network Admin role to the networking team, and assign the Compute Admin role to the development team.
C. 1. Create a project with a Shared VPC and assign the Network Admin role to the networking team.
2. Create a second project without a VPC, configure it as a Shared VPC service project, and assign the Compute Admin role to the development team.
D. 1. Create a project with a standalone VPC and assign the Network Admin role to the networking team.
2.Create a second project with a standalone VPC and assign the Compute Admin role to the development team.

3.Use VPC Peering to join the two VPCs.

Answer: (SHOW ANSWER)

Explanation

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

https://cloud.google.com/iam/docs/job-functions/networking#single_team_manages_security_network_for_organ Reference: https://cloud.google.com/vpc/docs/shared-vpc

NEW QUESTION: 39

You need to ensure reliability for your application and operations by supporting reliable task a scheduling for compute on GCP. Leveraging Google best practices, what should you do?

A. Using the Cron service provided by App Engine, publishing messages directly to a message-processing utility service running on Compute Engine instances.

B. Using the Cron service provided by App Engine, publish messages to a Cloud Pub/Sub topic. Subscribe to that topic using a message-processing utility service running on Compute Engine instances.

C. Using the Cron service provided by Google Kubernetes Engine (GKE), publish messages directly to a message-processing utility service running on Compute Engine instances.

D. Using the Cron service provided by GKE, publish messages to a Cloud Pub/Sub topic. Subscribe to that topic using a message-processing utility service running on Compute Engine instances.

Answer: B (LEAVE A REPLY)

Explanation

https://cloud.google.com/solutions/reliable-task-scheduling-compute-engine

NEW QUESTION: 40

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

A. Store time-series data from the game servers in Google Bigtable, and view it using Google Data Studio.

B. Output custom metrics to Stackdriver from the game servers, and create a Dashboard in Stackdriver Monitoring Console to view them.

C. Schedule BigQuery load jobs to ingest analytics files uploaded to Cloud Storage every ten minutes, and visualize the results in Google Data Studio.

D. Insert the KPIs into Cloud Datastore entities, and run ad hoc analysis and visualizations of them in Cloud Datalab.

Answer: A (LEAVE A REPLY)

Explanation

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

NEW QUESTION: 41

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

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

A. Assign custom IAM roles to the Google Groups you created in order to enforce security requirements. Encrypt data with a customer-supplied encryption key when storing files in Cloud Storage.

B. Assign custom IAM roles to the Google Groups you created in order to enforce security requirements. Enable default storage encryption before storing files in Cloud Storage.

C. Assign predefined IAM roles to the Google Groups you created in order to enforce security requirements.

Utilize Google's default encryption at rest when storing files in Cloud Storage.

D. Assign predefined IAM roles to the Google Groups you created in order to enforce security requirements. Ensure that the default Cloud KMS key is set before storing files in Cloud Storage.

Answer: D (LEAVE A REPLY)

Explanation

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

Topic 6, 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: 42

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

A. Use a different database.

B. Choose larger instances for your database.

C. Create snapshots of your database more regularly.

D. Implement routinely scheduled failovers of your databases.

Answer: D (LEAVE A REPLY)

Explanation

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

NEW QUESTION: 43

Your company wants to track whether someone is present in a meeting room reserved for a scheduled meeting.

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

A. Flat file

B. NoSQL

C. Relational

D. Blobstore

Answer: B (LEAVE A REPLY)

Explanation

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

NoSQL fits well for:

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

NEW QUESTION: 44

For this question, refer to the Mountkirk Games case study.
Mountkirk Games wants to set up a continuous delivery pipeline. Their architecture includes many small services that they want to be able to update and roll back quickly. Mountkirk Games has the following requirements:
* Services are deployed redundantly across multiple regions in the US and Europe.
* Only frontend services are exposed on the public internet.
* They can provide a single frontend IP for their fleet of services.
* Deployment artifacts are immutable.
Which set of products should they use?
A. Google Cloud Storage, Google App Engine, Google Network Load Balancer
B. Google Kubernetes Registry, Google Container Engine, Google HTTP(S) Load Balancer
C. Google Cloud Functions, Google Cloud Pub/Sub, Google Cloud Deployment Manager
D. Google Cloud Storage, Google Cloud Dataflow, Google Compute Engine
Answer: B (LEAVE A REPLY)

NEW QUESTION: 45
Your company is migrating its on-premises data center into the cloud. As part of the migration, you want to integrate Kubernetes Engine for workload orchestration. Parts of your architecture must also be PCI DSScompliant.
Which of the following is most accurate?
A. App Engine is the only compute platform on GCP that is certified for PCI DSS hosting.
B. Kubernetes Engine cannot be used under PCI DSS because it is considered shared hosting.
C. Kubernetes Engine and GCP provide the tools you need to build a PCI DSS-compliant environment.
D. All Google Cloud services are usable because Google Cloud Platform is certified PCI-compliant.
Answer: (SHOW ANSWER)
Explanation
https://cloud.google.com/security/compliance/pci-dss

NEW QUESTION: 46
You want to make a copy of a production Linux virtual machine in the US-Central region. You want to manage and replace the copy easily if there are changes on the production virtual machine. You will deploy the copy as a new instances in a different project in the US-East region. What steps must you take?
A. Use the Linux dd and netcat command to copy and stream the root disk contents to a new virtual machine instance in the US-East region.
B. Create a snapshot of the root disk and select the snapshot as the root disk when you create a new virtual machine instance in the US-East region.
C. Create an image file from the root disk with Linux dd command, create a new disk from the image file, and use it to create a new virtual machine instance in the US-East region
D. Create a snapshot of the root disk, create an image file in Google Cloud Storage from the snapshot, and create a new virtual machine instance in the US-East region using the image file for the root disk.

Answer: D (LEAVE A REPLY)
Explanation
https://stackoverflow.com/questions/36441423/migrate-google-compute-engine-instance-to-a-different-region

NEW QUESTION: 47
You team needs to create a Google Kubernetes Engine (GKE) cluster to host a newly built application that requires access to third-party services on the internet. Your company does not allow any Compute Engine instance to have a public IP address on Google Cloud. You need to create a deployment strategy that adheres to these guidelines. What should you do?
A. Create a Compute Engine instance, and install a NAT Proxy on the instance. Configure all workloads on GKE to pass through this proxy to access third-party services on the Internet
B. Configure the GKE cluster as a private cluster, and configure Cloud NAT Gateway for the cluster subnet
C. Configure the GKE cluster as a route-based cluster. Configure Private Google Access on the Virtual Private Cloud (VPC)
D. Configure the GKE cluster as a private cluster. Configure Private Google Access on the Virtual Private Cloud (VPC)
Answer: (SHOW ANSWER)
Explanation
A Cloud NAT gateway can perform NAT for nodes and Pods in a private cluster, which is a type of VPC-native cluster. The Cloud NAT gateway must be configured to apply to at least the following subnet IP address ranges for the subnet that your cluster uses:
Subnet primary IP address range (used by nodes)
Subnet secondary IP address range used for Pods in the cluster
Subnet secondary IP address range used for Services in the cluster
The simplest way to provide NAT for an entire private cluster is to configure a Cloud NAT gateway to apply to all of the cluster's subnet's IP address ranges.
https://cloud.google.com/nat/docs/overview

NEW QUESTION: 48

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: B (LEAVE A REPLY)

Explanation

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: 49

For this question, refer to the Dress4Win case study.

Dress4Win has end-to-end tests covering 100% of their endpoints. They want to ensure that the move to the cloud does not introduce any new bugs. Which additional testing methods should the developers employ to prevent an outage?

A. They should enable Google Stackdriver Debugger on the application code to show errors in the code.

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

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

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

Answer: B (LEAVE A REPLY)

Topic 5, Dress4Win Case 2

Company Overview

Dress4win is a web-based company that helps their users organize and manage their personal wardrobe using a website and mobile application. The company also cultivates an active social network that connects their users with designers and retailers. They monetize their services through advertising, e-commerce, referrals, and a freemium app model. The application has grown from a few servers in the founder's garage to several hundred servers and appliances in a collocated data center. However, the capacity of their infrastructure is now insufficient for the application's rapid growth. Because of this

growth and the company's desire to innovate faster. Dress4Win is committing to a full migration to a public cloud.

Solution Concept

For the first phase of their migration to the cloud, Dress4win is moving their development and test environments. They are also building a disaster recovery site, because their current infrastructure is at a single location. They are not sure which components of their architecture they can migrate as is and which components they need to change before migrating them.

Existing Technical Environment

The Dress4win application is served out of a single data center location. All servers run Ubuntu LTS v16.04.

Databases:

MySQL. 1 server for user data, inventory, static data:

- MySQL 5.8
- 8 core CPUs
- 128 GB of RAM
- 2x 5 TB HDD (RAID 1)

Redis 3 server cluster for metadata, social graph, caching. Each server is:

- Redis 3.2
- 4 core CPUs
- 32GB of RAM

Compute:

40 Web Application servers providing micro-services based APIs and static content.

- Tomcat - Java
- Nginx
- 4 core CPUs
- 32 GB of RAM

20 Apache Hadoop/Spark servers:

- Data analysis
- Real-time trending calculations
- 8 core CPUS
- 128 GB of RAM
- 4x 5 TB HDD (RAID 1)

3 RabbitMQ servers for messaging, social notifications, and events:

- 8 core CPUs
- 32GB of RAM

Miscellaneous servers:

- Jenkins, monitoring, bastion hosts, security scanners
- 8 core CPUs
- 32GB of RAM

Storage appliances:

iSCSI for VM hosts

Fiber channel SAN - MySQL databases
- 1 PB total storage; 400 TB available
NAS - image storage, logs, backups
- 100 TB total storage; 35 TB available
Business Requirements
Build a reliable and reproducible environment with scaled parity of production.
Improve security by defining and adhering to a set of security and Identity and Access Management (IAM) best practices for cloud.
Improve business agility and speed of innovation through rapid provisioning of new resources.
Analyze and optimize architecture for performance in the cloud.
Technical Requirements
Easily create non-production environment in the cloud.
Implement an automation framework for provisioning resources in cloud.
Implement a continuous deployment process for deploying applications to the on-premises datacenter or cloud.
Support failover of the production environment to cloud during an emergency.
Encrypt data on the wire and at rest.
Support multiple private connections between the production data center and cloud environment.
Executive Statement
Our investors are concerned about our ability to scale and contain costs with our current infrastructure. They are also concerned that a competitor could use a public cloud platform to offset their up-front investment and free them to focus on developing better features. Our traffic patterns are highest in the mornings and weekend evenings; during other times, 80% of our capacity is sitting idle.
Our capital expenditure is now exceeding our quarterly projections. Migrating to the cloud will likely cause an initial increase in spending, but we expect to fully transition before our next hardware refresh cycle. Our total cost of ownership (TCO) analysis over the next 5 years for a public cloud strategy achieves a cost reduction between 30% and 50% over our current model.

NEW QUESTION: 50
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 smaller changes to production.
B. Deploy changes to a small subset of users before rolling out to production.
C. Deploy fewer changes to production.
D. Increase the load on your test and staging environments.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 51

You need to design a solution for global load balancing based on the URL path being requested. You need to ensure operations reliability and end-to-end in-transit encryption based on Google best practices.
What should you do?
A. Create a cross-region load balancer with URL Maps.
B. Create an HTTPS load balancer with URL maps.
C. Create appropriate instance groups and instances. Configure SSL proxy load balancing.
D. Create a global forwarding rule. Configure SSL proxy balancing.
Answer: B (LEAVE A REPLY)
Explanation
Reference https://cloud.google.com/load-balancing/docs/https/url-map

NEW QUESTION: 52
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. 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.
B. 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.
C. Use a load testing tool to simulate the expected number of concurrent users and total requests to your application, and inspect the results.
D. Use Cloud Debugger in the development environment to understand the latency between the different microservices.
Answer: (SHOW ANSWER)

NEW QUESTION: 53
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)
Explanation
https://cloud.google.com/bigquery/docs/managing-partitioned-tables

NEW QUESTION: 54

Your company uses the Firewall Insights feature in the Google Network Intelligence Center. You have several firewall rules applied to Compute Engine instances. You need to evaluate the efficiency of the applied firewall ruleset. When you bring up the Firewall Insights page in the Google Cloud Console, you notice that there are no log rows to display. What should you do to troubleshoot the issue?
A. Enable Virtual Private Cloud (VPC) flow logging.
B. Enable Firewall Rules Logging for the firewall rules you want to monitor.
C. Verify that your user account is assigned the compute.networkAdmin Identity and Access Management (IAM) role.
D. Install the Google Cloud SDK, and verify that there are no Firewall logs in the command line output.
Answer: B (LEAVE A REPLY)
Reference: https://cloud.google.com/network-intelligence-center/docs/firewall-insights/how-to/using-firewall- insights

NEW QUESTION: 55
You want to automate the creation of a managed instance group and a startup script to install the OS package dependencies. You want to minimize the startup time for VMs in the instance group.
What should you do?
A. Use Terraform to create the managed instance group and a startup script to install the OS package dependencies.
B. Create a custom VM image with all OS package dependencies. Use Deployment Manager to create the managed instance group with the VM image.
C. Use Puppet to create the managed instance group and install the OS package dependencies.
D. Use Deployment Manager to create the managed instance group and Ansible to install the OS package dependencies.
Answer: B (LEAVE A REPLY)
Explanation
"Custom images are more deterministic and start more quickly than instances with startup scripts. However, startup scripts are more flexible and let you update the apps and settings in your instances more easily."
https://cloud.google.com/compute/docs/instance-templates/create-instance-templates#using_custom_or_public_i

NEW QUESTION: 56
Your operations team currently stores 10 TB of data m an object storage service from a third-party provider.
They want to move this data to a Cloud Storage bucket as quickly as possible, following Google-recommended practices. They want to minimize the cost of this data migration. When approach should they use?
A. Use the gsutil mv command lo move the data
B. Use the Storage Transfer Service to move the data
C. Download the data to a Transfer Appliance and ship it to Google

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

Answer: (SHOW ANSWER)

Explanation

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

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

NEW QUESTION: 57

For this question, refer to the Dress4Win case study.

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

Which configuration should Dress4Win use?

A. Use a distributed file system to store customers' images. As storage needs increase, add more persistent disks and/or nodes. Use a Google Cloud SQL database to maintain metadata that maps each customer's ID to their image files.

B. Store image files in a Google Cloud Storage bucket. Add custom metadata to the uploaded images in Cloud Storage that contains the customer's unique ID.

C. Store image files in a Google Cloud Storage bucket. Use Google Cloud Datastore to maintain metadata that maps each customer's ID and their image files.

D. Use a distributed file system to store customers' images. As storage needs increase, add more persistent disks and/or nodes. Assign each customer a unique ID, which sets each file's owner attribute, ensuring privacy of images.

Answer: C (LEAVE A REPLY)

NEW QUESTION: 58

You have an application that makes HTTP requests to Cloud Storage. Occasionally the requests fail with HTTP status codes of 5xx and 429.

How should you handle these types of errors?

A. Use gRPC instead of HTTP for better performance.

B. Implement retry logic using a truncated exponential backoff strategy.

C. Make sure the Cloud Storage bucket is multi-regional for geo-redundancy.

D. Monitor https://status.cloud.google.com/feed.atom and only make requests if Cloud Storage is not reporting an incident.

Answer: A (LEAVE A REPLY)

Explanation

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

NEW QUESTION: 59

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

NEW QUESTION: 60
Your company wants to start using Google Cloud resources but wants to retain their on-premises Active Directory domain controller for identity management. What should you do?
A. Use the Admin Directory API to authenticate against the Active Directory domain controller.
B. Use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML SSO.
C. Use Cloud Identity-Aware Proxy configured to use the on-premises Active Directory domain controller as an identity provider.
D. Use Compute Engine to create an Active Directory (AD) domain controller that is a replica of the onpremises AD domain controller using Google Cloud Directory Sync.
Answer: B (LEAVE A REPLY)
Explanation
https://cloud.google.com/solutions/federating-gcp-with-active-directory-introduction#implementing_federation

NEW QUESTION: 61
Your company pushes batches of sensitive transaction data from its application server VMs to Cloud Pub/Sub for processing and storage. What is the Google-recommended way for your application to authenticate to the required Google Cloud services?
A. Ensure that VM service accounts are granted the appropriate Cloud Pub/Sub IAM roles.
B. Ensure that VM service accounts do not have access to Cloud Pub/Sub, and use VM access scopes to grant the appropriate Cloud Pub/Sub IAM roles.
C. Generate an OAuth2 access token for accessing Cloud Pub/Sub, encrypt it, and store it in Cloud Storage for access from each VM.
D. Create a gateway to Cloud Pub/Sub using a Cloud Function, and grant the Cloud Function service account the appropriate Cloud Pub/Sub IAM roles.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 62
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 process can be automated with Migrate for Compute Engine.
B. The new approach will be significantly less costly, make it easier to manage the underlying infrastructure, and automatically manage the CI/CD pipelines.
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 monolithic solution can be converted to a container with Docker. The generated container can then be deployed into a Kubernetes cluster.
Answer: (SHOW ANSWER)
Explanation
The new approach will make it easier to decouple infrastructure from an application, develop and release new features, manage the underlying infrastructure, manage CI/CD pipelines and perform A/B testing, and scale the solution if necessary.

NEW QUESTION: 63
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)
Explanation
https://cloud.google.com/solutions/gaming/cloud-game-infrastructure#dedicated_game_server

NEW QUESTION: 64
For this question, refer to the Helicopter Racing League (HRL) case study. HRL is looking for a cost-effective approach for storing their race data such as telemetry. They want to keep all historical records, train models using only the previous season's data, and plan for data growth in terms of volume and information collected.
You need to propose a data solution. Considering HRL business requirements and the goals expressed by CEO S. Hawke, what should you do?
A. Use Firestore for its scalable and flexible document-based database. Use collections to aggregate race data by season and event.
B. Use Cloud Spanner for its scalability and ability to version schemas with zero downtime. Split race data using season as a primary key.
C. Use BigQuery for its scalability and ability to add columns to a schema. Partition race data based on season.
D. Use Cloud SQL for its ability to automatically manage storage increases and compatibility with MySQL. Use separate database instances for each season.
Answer: (SHOW ANSWER)
Reference: https://cloud.google.com/bigquery/public-data

NEW QUESTION: 65
For this question, refer to the TerramEarth case study. You are building a microservice-based application for TerramEarth. The application is based on Docker containers. You want to follow Google-recommended practices to build the application continuously and store the build artifacts. What should you do?
A. 1. Configure a trigger in Cloud Build for new source changes.
2. Invoke Cloud Build to build container images for each microservice, and tag them using the code commit hash.
3. Push the images to the Artifact Registry.
B. 1 Create a Scheduler job to check the repo every minute.
2. For any new change, invoke Cloud Build to build container images for the microservices.
3. Tag the images using the current timestamp, and push them to the Artifact Registry.
C. 1. Configure a trigger in Cloud Build for new source changes.
2. Invoke Cloud Build to build one container image, and tag the image with the label 'latest.'
3. Push the image to the Artifact Registry.
D. 1. Configure a trigger in Cloud Build for new source changes.
2. The trigger invokes build jobs and build container images for the microservices.
3. Tag the images with a version number, and push them to Cloud Storage.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 66
You need to optimize batch file transfers into Cloud Storage for Mountkirk Games' new Google Cloud solution.
The batch files contain game statistics that need to be staged in Cloud Storage and be processed by an extract transform load (ETL) tool. What should you do?
A. Use gsutil to batch move files in sequence.
B. Use gsutil to batch copy the files in parallel.
C. Use gsutil to extract the files as the first part of ETL.
D. Use gsutil to load the files as the last part of ETL.
Answer: (SHOW ANSWER)
Reference: https://cloud.google.com/storage/docs/gsutil/commands/cp

NEW QUESTION: 67
You are implementing a single Cloud SQL MySQL second-generation database that contains business-critical transaction data. You want to ensure that the minimum amount of data is lost in case of catastrophic failure.
Which two features should you implement? (Choose two.)
A. Sharding
B. Read replicas
C. Binary logging
D. Automated backups
E. Semisynchronous replication
Answer: (SHOW ANSWER)
Explanation
Backups help you restore lost data to your Cloud SQL instance. Additionally, if an instance is having a problem, you can restore it to a previous state by using the backup to overwrite it. Enable automated backups for any instance that contains necessary data. Backups protect your data from loss or damage. Enabling automated backups, along with binary logging, is also required for some operations, such as clone and replica creation.
Reference: https://cloud.google.com/sql/docs/mysql/backup-recovery/backups

NEW QUESTION: 68
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. Migrate Nginx and Tomcat to App Engine. Deploy a MySQL server using Cloud Launcher. Deploy Jenkins to Compute Engine using Cloud Launcher.
C. Deploy Nginx and Tomcat using Cloud Launcher. Deploy a MySQL server using Cloud Launcher. Deploy Jenkins to Compute Engine using Cloud Deployment Manager scripts.

D. 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.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 69
Your company is moving 75 TB of data into Google Cloud. You want to use Cloud Storage and follow Googlerecommended practices. What should you do?
A. Move your data onto a Transfer Appliance. Use a Transfer Appliance Rehydrator to decrypt the data into Cloud Storage.
B. Move your data onto a Transfer Appliance. Use Cloud Dataprep to decrypt the data into Cloud Storage.
C. Install gsutil on each server that contains data. Use resumable transfers to upload the data into Cloud Storage.
D. Install gsutil on each server containing data. Use streaming transfers to upload the data into Cloud Storage.
Answer: A (LEAVE A REPLY)
Explanation
https://cloud.google.com/transfer-appliance/docs/2.0/faq

NEW QUESTION: 70
Your company has an enterprise application running on Compute Engine that requires high availability and high performance. The application has been deployed on two instances in two zones in the same region m active passive mode. The application writes data to a persistent disk in the case of a single zone outage that data should be immediately made available to the other instance in the other zone. You want to maximize performance while minimizing downtime and data loss. What should you do?
A. 1. Attach a regional SSD persistent Ask to the first instance
2. In case of a zone outage, force-attach the disk to the other 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 local SSD lo the first instance disk
2. Execute an rsync command every hour where the target is a persistent SSD disk attached to the second instance
3. In case of a zone outage, use the second instance
D. 1. Attach a 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
Answer: A (LEAVE A REPLY)

NEW QUESTION: 71

You are using Cloud CDN to deliver static HTTP(S) website content hosted on a Compute Engine instance group. You want to improve the cache hit ratio.
What should you do?
A. Customize the cache keys to omit the protocol from the key.
B. Shorten the expiration time of the cached objects.
C. Make sure the HTTP(S) header "Cache-Region" points to the closest region of your users.
D. Replicate the static content in a Cloud Storage bucket. Point CloudCDN toward a load balancer on that bucket.
Answer: (SHOW ANSWER)
Explanation
Reference https://cloud.google.com/cdn/docs/bestpractices#
using_custom_cache_keys_to_improve_cache_hit_ratio

NEW QUESTION: 72
Dress4win has end to end tests covering 100% of their endpoints.
They want to ensure that the move of cloud does not introduce any new bugs.
Which additional testing methods should the developers employ to prevent an outage?
A. They should run the end to end tests in the cloud staging environment to determine if the code is working as intended.
B. They should add additional unit tests and production scale load tests on their cloud staging environment.
C. They should enable google stack driver debugger on the application code to show errors in the code
D. They should add canary tests so developers can measure how much of an impact the new release causes to latency
Answer: (SHOW ANSWER)

NEW QUESTION: 73
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. Create a Cloud Function that reads data from BigQuery and cleans it. Trigger it. Trigger the Cloud Function from a Compute Engine instance.
B. 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.
C. Use Cloud Dataprep and configure the BigQuery tables as the source. Schedule a daily job to clean the data.
D. Set up a streaming Cloud Dataflow job, receiving data by the ingestion process. Clean the data in a Cloud Dataflow pipeline.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 74
Your company is building a new architecture to support its data-centric business focus. You are responsible for setting up the network. Your company's mobile and web-facing applications will be deployed on-premises, and all data analysis will be conducted in GCP. The plan is to process and load 7 years of archived .csv files totaling 900 TB of data and then continue loading 10 TB of data daily. You currently have an existing 100-MB internet connection.
What actions will meet your company's needs?
A. Compress and upload both achieved files and files uploaded daily using the qsutil -m option.
B. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transfer archived data to Cloud Storage. Establish a connection with Google using a Dedicated Interconnect or Direct Peering connection and use it to upload files daily.
C. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transfer archived data to Cloud Storage. Establish one Cloud VPN Tunnel to VPC networks over the public internet, and compares and upload files daily using the gsutil -m option.
D. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storage. Establish a Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily.
Answer: B (LEAVE A REPLY)
Explanation
https://cloud.google.com/interconnect/docs/how-to/direct-peering

NEW QUESTION: 75
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. Validate that the Service Account used by the Cloud SQL proxy container still has the Cloud Build Editor role.
C. Use gcloud sql instances restart.
D. In the GCP Console, navigate to Cloud SQL. Restore the latest backup. Use kubect1 to restart all pods.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 76
Your company has an application running on multiple Compute Engine instances. You need to ensure that the application can communicate with an on-premises service that requires high throughput via internal IPs, while minimizing latency. What should you do?
A. Use OpenVPN to configure a VPN tunnel between the on-premises environment and Google Cloud.
B. Configure a direct peering connection between the on-premises environment and Google Cloud.

C. Use Cloud VPN to configure a VPN tunnel between the on-premises environment and Google Cloud.
D. Configure a Cloud Dedicated Interconnect connection between the on-premises environment and Google Cloud.
Answer: D (LEAVE A REPLY)
Explanation
Reference https://cloud.google.com/architecture/setting-up-private-access-to-cloud-apis-through-vpn-tunnels

NEW QUESTION: 77
For this question, refer to the Mountkirk Games case study. Which managed storage option meets Mountkirk's technical requirement for storing game activity in a time series database service?
A. Cloud Bigtable
B. Cloud Spanner
C. BigQuery
D. Cloud Datastore
Answer: A (LEAVE A REPLY)
Explanation
https://cloud.google.com/blog/products/databases/getting-started-with-time-series-trend-predictions-using-gcp

NEW QUESTION: 78
Your company captures all web traffic data in Google Analytics 260 and stores it in BigQuery. Each country has its own dataset. Each dataset has multiple tables. You want analysts from each country to be able to see and query only the data for their respective countries.
How should you configure the access rights?
A. Create a group per country. Add analysts to their respective country-groups. Create a single group 'all_analysts', and add all country-groups as members. Grant the 'all-analysis' group the IAM role of BigQuery jobUser. Share the appropriate dataset with view access with each respective analyst country-group.
B. Create a group per country. Add analysts to their respective country-groups. Create a single group

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

NEW QUESTION: 79
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)
Explanation
https://cloud.google.com/solutions/data-lifecycle-cloud-platform
https://cloud.google.com/solutions/designing-connected-vehicle-platform

NEW QUESTION: 80
You are moving an application that uses MySQL from on-premises to Google Cloud. The application will run on Compute Engine and will use Cloud SQL. You want to cut over to the Compute Engine deployment of the application with minimal downtime and no data loss to your customers. You want to migrate the application with minimal modification. You also need to determine the cutover strategy. What should you do?
A. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server.
2.Stop the on-premises application.
3.Create a mysqldump of the on-premises MySQL server.
4.Upload the dump to a Cloud Storage bucket.
5.Import the dump into Cloud SQL.
6.Modify the source code of the application to write queries to both databases and read from its local database.
7.Start the Compute Engine application.

8.Stop the on-premises application.
B. 1. Set up Cloud SQL proxy and MySQL proxy.
2.Create a mysqldump of the on-premises MySQL server.
3.Upload the dump to a Cloud Storage bucket.
4.Import the dump into Cloud SQL.
5.Stop the on-premises application.
6.Start the Compute Engine application.
C. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server.
2.Stop the on-premises application.
3.Start the Compute Engine application, configured to read and write to the on-premises MySQL server.
4.Create the replication configuration in Cloud SQL.
5.Configure the source database server to accept connections from the Cloud SQL replica.
6.Finalize the Cloud SQL replica configuration.
7.When replication has been completed, stop the Compute Engine application.
8.Promote the Cloud SQL replica to a standalone instance.
9.Restart the Compute Engine application, configured to read and write to the Cloud SQL standalone instance.
D. 1. Stop the on-premises application.
2.Create a mysqldump of the on-premises MySQL server.
3.Upload the dump to a Cloud Storage bucket.
4.Import the dump into Cloud SQL.
5.Start the application on Compute Engine.
Answer: (SHOW ANSWER)
Explanation
External replica promotion migration In the migration strategy of external replica promotion, you create an external database replica and synchronize the existing data to that replica. This can happen with minimal downtime to the existing database. When you have a replica database, the two databases have different roles that are referred to in this document as primary and replica. After the data is synchronized, you promote the replica to be the primary in order to move the management layer with minimal impact to database uptime. In Cloud SQL, an easy way to accomplish the external replica promotion is to use the automated migration workflow. This process automates many of the steps that are needed for this type of migration.
https://cloud.google.com/architecture/migrating-mysql-to-cloudsql-concept
- The best option for migrating your MySQL database is to use an external replica promotion. In this strategy, you create a replica database and set your existing database as the primary. You wait until the two databases are in sync, and you then promote your MySQL replica database to be the primary. This process minimizes database downtime related to the database migration. -
https://cloud.google.com/architecture/migrating-mysql-to-cloudsql-concept#external_replica_promotion_migrati

NEW QUESTION: 81
You are migrating your on-premises solution to Google Cloud in several phases. You will use Cloud VPN to maintain a connection between your on-premises systems and Google Cloud until the migration is completed.
You want to make sure all your on-premises systems remain reachable during this period. How should you organize your networking in Google Cloud?
A. Use the same IP range on Google Cloud as you use on-premises
B. Use an IP range on Google Cloud that does not overlap with the range you use on-premises for your primary IP range and use a secondary range with the same IP range as you use on-premises
C. Use the same IP range on Google Cloud as you use on-premises for your primary IP range and use a secondary range that does not overlap with the range you use on-premises
D. Use an IP range on Google Cloud that does not overlap with the range you use on-premises
Answer: D (LEAVE A REPLY)

NEW QUESTION: 82
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)
Explanation
For all Google Cloud services secured with VPC Service Controls, you can ensure that: Resources within a perimeter are accessed only from clients within authorized VPC networks using Private Google Access with either Google Cloud or on-premises. https://cloud.google.com/vpc-service-controls/docs/overview
https://cloud.google.com/vpc-service-controls/docs/overview. You create a service control across your VPC and any cloud bucket or any project resource to restrict access. Anything outside of it can't access the resources within service control perimeter

NEW QUESTION: 83
For this question, refer to the EHR Healthcare case study. You are a developer on the EHR customer portal team. Your team recently migrated the customer portal application to Google Cloud. The load has

increased on the application servers, and now the application is logging many timeout errors. You recently incorporated Pub/Sub into the application architecture, and the application is not logging any Pub/Sub publishing errors.
You want to improve publishing latency. What should you do?
A. Increase the Pub/Sub Total Timeout retry value.
B. Move from a Pub/Sub subscriber pull model to a push model.
C. Turn off Pub/Sub message batching.
D. Create a backup Pub/Sub message queue.
Answer: C (LEAVE A REPLY)
Explanation
https://cloud.google.com/pubsub/docs/publisher?hl=en#batching

NEW QUESTION: 84
You are creating a solution to remove backup files older than 90 days from your backup Cloud Storage bucket.
You want to optimize ongoing Cloud Storage spend. What should you do?
A. Write a lifecycle management rule in XML and push it to the bucket with gsutil.
B. Write a lifecycle management rule in JSON and push it to the bucket with gsutil.
C. Schedule a cron script using gsutil is -lr gs://backups/** to find and remove items older than 90 days.
D. Schedule a cron script using gsutil ls -1 gs://backups/** to find and remove items older than 90 days and schedule it with cron.
Answer: B (LEAVE A REPLY)
Explanation
https://cloud.google.com/storage/docs/gsutil/commands/lifecycle

NEW QUESTION: 85
You want to enable your running Google Kubernetes Engine cluster to scale as demand for your application changes.
What should you do?
A. Add additional nodes to your Kubernetes Engine cluster using the following command:gcloud container clusters resizeCLUSTER_Name - -size 10
B. Add a tag to the instances in the cluster with the following command:gcloud compute instances add-tagsINSTANCE - -tags enable-autoscaling max-nodes-10
C. Update the existing Kubernetes Engine cluster with the following command:gcloud alpha container clustersupdate mycluster - -enable-autoscaling - -min-nodes=1 - -max-nodes=10
D. Create a new Kubernetes Engine cluster with the following command:gcloud alpha container clusterscreate mycluster - -enable-autoscaling - -min-nodes=1 - -max-nodes=10and redeploy your application
Answer: C (LEAVE A REPLY)
Explanation

https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler To enable autoscaling for an existing node pool, run the following command:
gcloud container clusters update [CLUSTER_NAME] --enable-autoscaling \--min-nodes 1 --max-nodes 10
--zone [COMPUTE_ZONE] --node-pool default-pool

NEW QUESTION: 86
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)
Explanation
The speed of deployment can be changed by limiting the size of the uploaded app, limiting the complexity of the build necessary in the Dockerfile, if present, and by ensuring a fast and reliable internet connection.
Note: Alpine Linux is built around musl libc and busybox. This makes it smaller and more resource efficient than traditional GNU/Linux distributions. A container requires no more than 8 MB and a minimal installation to disk requires around 130 MB of storage. Not only do you get a fully-fledged Linux environment but a large selection of packages from the repository.
References: https://groups.google.com/forum/#!topic/google-appengine/hZMEkmmObDU
https://www.alpinelinux.org/about/

NEW QUESTION: 87
For this question, refer to the EHR Healthcare case study. You need to define the technical architecture for hybrid connectivity between EHR's on-premises systems and Google Cloud. You want to follow Google's recommended practices for production-level applications. Considering the EHR Healthcare business and technical requirements, what should you do?
A. Configure two Partner Interconnect connections in one metro (City), and make sure the Interconnect connections are placed in different metro zones.

B. Configure two VPN connections from on-premises to Google Cloud, and make sure the VPN devices on-premises are in separate racks.
C. Configure Direct Peering between EHR Healthcare and Google Cloud, and make sure you are peering at least two Google locations.
D. Configure two Dedicated Interconnect connections in one metro (City) and two connections in another metro, and make sure the Interconnect connections are placed in different metro zones.
Answer: D (LEAVE A REPLY)
Explanation
based on the requirement of secure and high-performance connection between on-premises systems to Google Cloud
https://cloud.google.com/network-connectivity/docs/interconnect/tutorials/partner-creating-9999-availability

NEW QUESTION: 88
For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the database workloads for your company, Mountkirk Games. Considering the business and technical requirements, what should you do?
A. Use Cloud SQL for time series data, and use Cloud Bigtable for historical data queries.
B. Use Cloud SQL to replace MySQL, and use Cloud Spanner for historical data queries.
C. Use Cloud Bigtable to replace MySQL, and use BigQuery for historical data queries.
D. Use Cloud Bigtable for time series data, use Cloud Spanner for transactional data, and use BigQuery for historical data queries.
Answer: (SHOW ANSWER)
Explanation
https://cloud.google.com/bigtable/docs/schema-design-time-series

NEW QUESTION: 89
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. Log into each Compute Engine instance and collect disk, CPU, memory, and network usage statistics for analysis.
B. Use the gcloud compute instances list to list the virtual machine instances that have the idle: true label set.
C. Use the gcloud recommender command to list the idle virtual machine instances.
D. From the Google Console, identify which Compute Engine instances in the managed instance groups are no longer responding to health check probes.
Answer: (SHOW ANSWER)

Reference: https://cloud.google.com/compute/docs/instances/viewing-and-applying-idle-vm-recommendations

NEW QUESTION: 90
For this question, refer to the TerramEarth case study.
The TerramEarth development team wants to create an API to meet the company's business requirements. You want the development team to focus their development effort on business value versus creating a custom framework. Which method should they use?
A. Use Google App Engine with Google Cloud Endpoints. Focus on an API for dealers and partners.
B. Use Google App Engine with a JAX-RS Jersey Java-based framework. Focus on an API for the public.
C. Use Google App Engine with the Swagger (open API Specification) framework. Focus on an API for the public.
D. Use Google Container Engine with a Django Python container. Focus on an API for the public.
E. Use Google Container Engine with a Tomcat container with the Swagger (Open API Specification) framework. Focus on an API for dealers and partners.
Answer: A (LEAVE A REPLY)
Explanation
https://cloud.google.com/endpoints/docs/openapi/about-cloud-endpoints?hl=en_US&_ga=2.21787131.-1712523
https://cloud.google.com/endpoints/docs/openapi/architecture-overview
https://cloud.google.com/storage/docs/gsutil/commands/test
Develop, deploy, protect and monitor your APIs with Google Cloud Endpoints. Using an Open API Specification or one of our API frameworks, Cloud Endpoints gives you the tools you need for every phase of API development.
From scenario:
Business Requirements
Decrease unplanned vehicle downtime to less than 1 week, without increasing the cost of carrying surplus inventory Support the dealer network with more data on how their customers use their equipment to better position new products and services Have the ability to partner with different companies - especially with seed and fertilizer suppliers in the fast-growing agricultural business - to create compelling joint offerings for their customers.
Reference: https://cloud.google.com/certification/guides/cloud-architect/casestudy-terramearth

NEW QUESTION: 91
Your company is using BigQuery as its enterprise data warehouse. Data is distributed over several Google Cloud projects. All queries on BigQuery need to be billed on a single project. You want to make sure that no query costs are incurred on the projects that contain the data. Users should be able to query the datasets, but not edit them.
How should you configure users' access roles?

A. Add all users to a group. Grant the group the role of BigQuery user on the billing project and BigQuery dataViewer on the projects that contain the data.
B. Add all users to a group. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery user on the projects that contain the data.
C. Add all users to a group. Grant the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data.
D. Add all users to a group. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery jobUser on the projects that contain the data.
Answer: (SHOW ANSWER)
Reference: https://cloud.google.com/bigquery/docs/running-queries

NEW QUESTION: 92
Your company has a Google Cloud project that uses BigQuery for data warehousing on a pay-per-use basis.
You want to monitor queries in real time to discover the most costly queries and which users spend the most.
What should you do?
A. 1. Create a Cloud Logging sink to export BigQuery data access logs to Cloud Storage. 2. Develop a Dataflow pipeline to compute the cost of queries split by users.
B. 1. Create a Cloud Logging sink to export BigQuery data access logs to BigQuery.
2. Perform a BigQuery query on the generated table to extract the information you need.
C. 1. Activate billing export into BigQuery.
2. Perform a BigQuery query on the billing table to extract the information you need.
D. 1. In the BigQuery dataset that contains all the tables to be queried, add a label for each user that can launch a query.
2. Open the Billing page of the project.
3. Select Reports.
4. Select BigQuery as the product and filter by the user you want to check.
Answer: C (LEAVE A REPLY)
Explanation
https://cloud.google.com/blog/products/data-analytics/taking-a-practical-approach-to-bigquery-cost-monitoring

NEW QUESTION: 93
Your customer is moving their corporate applications to Google Cloud Platform. The security team wants detailed visibility of all projects in the organization. You provision the Google Cloud Resource Manager and set up yourself as the org admin. What Google Cloud Identity and Access Management (Cloud IAM) roles should you give to the security team'?
A. Org viewer, project owner
B. Org viewer, project viewer
C. Org admin, project browser
D. Project owner, network admin
Answer: B (LEAVE A REPLY)
Explanation
https://cloud.google.com/iam/docs/using-iam-securely

NEW QUESTION: 94
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: C (LEAVE A REPLY)
Explanation
https://cloud.google.com/bigquery/docs/managing-partitioned-tables#partition-expiration
https://cloud.google.com/storage/docs/lifecycle

NEW QUESTION: 95
Your customer support tool logs all email and chat conversations to Cloud Bigtable for retention and analysis.

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

A. Hash all data using SHA256

B. Encrypt all data using elliptic curve cryptography

C. De-identify the data with the Cloud Data Loss Prevention API

D. Use regular expressions to find and redact phone numbers, email addresses, and credit card numbers

Answer: A (LEAVE A REPLY)

Reference: https://cloud.google.com/solutions/pci-dss-compliance-ingcp#

NEW QUESTION: 96

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

A. Create a Kubernetes admission controller to prevent the container from starting if it is not approved for usage in the given environment

B. Configure a Kubernetes lifecycle hook to prevent the container from starting if it is not approved for usage in the given environment

C. Implement a corporate policy to prevent teams from deploying Docker image to an environment unless the Docker image was tested in an earlier environment

D. Configure the binary authorization policies for the development, staging and production clusters. Create attestations as part of the continuous integration pipeline"

Answer: D (LEAVE A REPLY)

Explanation

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

NEW QUESTION: 97

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

A. Web applications deployed using App Engine standard environment

B. Jenkins, monitoring, bastion hosts, security scanners services deployed on custom machine types

C. RabbitMQ deployed using an unmanaged instance group

D. Hadoop/Spark deployed using Cloud Dataproc Regional in High Availability mode

Answer: D (LEAVE A REPLY)

NEW QUESTION: 98
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)
Explanation
https://cloud.google.com/logging/docs/agent/default-logs

NEW QUESTION: 99
The operations manager asks you for a list of recommended practices that she should consider when migrating a J2EE application to the cloud. Which three practices should you recommend? Choose 3 answers
A. Port the application code to run on Google App Engine.
B. Integrate Cloud Dataflow into the application to capture real-time metrics.
C. Instrument the application with a monitoring tool like Stackdriver Debugger.
D. Select an automation framework to reliably provision the cloud infrastructure.
E. Deploy a continuous integration tool with automated testing in a staging environment.
F. Migrate from MySQL to a managed NoSQL database like Google Cloud Datastore or Bigtable.
Answer: A,E,F (LEAVE A REPLY)
Explanation
References: https://cloud.google.com/appengine/docs/standard/java/tools/uploadinganapp
https://cloud.google.com/appengine/docs/standard/java/building-app/cloud-sql

NEW QUESTION: 100
For this question, refer to the TerramEarth case study. You need to implement a reliable, scalable GCP solution for the data warehouse for your company, TerramEarth. Considering the TerramEarth business and technical requirements, what should you do?
A. Replace the existing data warehouse with BigQuery. Use table partitioning.
B. Replace the existing data warehouse with a Compute Engine instance with 96 CPUs.
C. Replace the existing data warehouse with BigQuery. Use federated data sources.
D. Replace the existing data warehouse with a Compute Engine instance with 96 CPUs. Add an additional Compute Engine pre-emptible instance with 32 CPUs.
Answer: (SHOW ANSWER)

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

NEW QUESTION: 101
You have deployed an application on Anthos clusters (formerly Anthos GKE). According to the SRE practices at your company you need to be alerted if the request latency is above a certain threshold for a specified amount of time. What should you do?
A. Enable the Cloud Trace API on your project and use Cloud Monitoring Alerts to send an alert based on the Cloud Trace metrics
B. Configure Anthos Config Management on your cluster and create a yaml file that defines the SLO and alerting policy you want to deploy in your cluster
C. Use Cloud Profiler to follow up the request latency. Create a custom metric in Cloud Monitoring based on the results of Cloud Profiler, and create an Alerting Policy in case this metric exceeds the threshold
D. Install Anthos Service Mesh on your cluster. Use the Google Cloud Console to define a Service Level Objective (SLO)
Answer: D (LEAVE A REPLY)
Explanation
https://cloud.google.com/service-mesh/docs/overview
https://cloud.google.com/service-mesh/docs/observability/slo-overview

NEW QUESTION: 102
For this question, refer to the Dress4Win case study.
Dress4Win has asked you for advice on how to migrate their on-premises MySQL deployment to the cloud.
They want to minimize downtime and performance impact to their on-premises solution during the migration.
Which approach should you recommend?
A. Create a dump of the on-premises MySQL master server, and then shut it down, upload it to the cloud environment, and load into a new MySQL cluster.
B. Setup a MySQL replica server/slave in the cloud environment, and configure it for asynchronous replication from the MySQL master server on-premises until cutover.
C. Create a new MySQL cluster in the cloud, configure applications to begin writing to both on-premises and cloud MySQL masters, and destroy the original cluster at cutover.
D. Create a dump of the MySQL replica server into the cloud environment, load it into: Google Cloud Datastore, and configure applications to read/write to Cloud Datastore at cutover.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 103
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 SSDbacked instances for VM boot/data volumes. Cloud Storage for log archives and thumbnails.
Answer: D (LEAVE A REPLY)
Explanation
https://cloud.google.com/compute/docs/disks

NEW QUESTION: 104
Auditors visit your teams every 12 months and ask to review all the Google Cloud Identity and Access Management (Cloud IAM) policy changes in the previous 12 months. You want to streamline and expedite the analysis and audit process. What should you do?
A. Create custom Google Stackdriver alerts and send them to the auditor.
B. Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor.
C. Use cloud functions to transfer log entries to Google Cloud SQL and use ACLS and views to limit an auditor's view.
D. Enable Google Cloud Storage (GCS) log export to audit logs Into a GCS bucket and delegate access to the bucket.
Answer: D (LEAVE A REPLY)
Explanation
Export the logs to Google Cloud Storage bucket - Archive Storage, as it will not be used for 1 year, price for which is $0.004 per GB per Month. The price for long term storage in BigQuery is $0.01 per GB per Month (250% more). Also for analysis purpose, whenever Auditors are there(once per year), you can use BigQuery and use GCS bucket as external data source. BigQuery supports querying Cloud Storage data from these storage classes:
Standard Nearline Coldline Archive

NEW QUESTION: 105
A news teed web service has the following code running on Google App Engine. During peak load, users report that they can see news articles they already viewed. What is the most likely cause of this problem?

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

app = Flask(_name_)
sessions = {}

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

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

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

     return news.render(news_articles)

if _name_ == "_main_":
     app.run()
```


A. The session variable is local to just a single instance.
B. The session variable is being overwritten in Cloud Datastore.
C. The URL of the API needs to be modified to prevent caching.
D. The HTTP Expires header needs to be set to -1 to stop caching.
Answer: (SHOW ANSWER)
Explanation
https://stackoverflow.com/questions/3164280/google-app-engine-cache-list-in-session-variable?rq=1

NEW QUESTION: 106
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)
Explanation
https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy
"Roles are always inherited, and there is no way to explicitly remove a permission for a lower-level resource that is granted at a higher level in the resource hierarchy. Given the above example, even if you were to remove the Project Editor role from Bob on the "Test GCP Project", he would still inherit that role from the
"Dept Y" folder, so he would still have the permissions for that role on "Test GCP Project"." Reference: https://cloud.google.com/resource-manager/docs/creating-managing-folders

NEW QUESTION: 107
You deploy your custom java application to google app engine.
It fails to deploy and gives you the following stack trace:

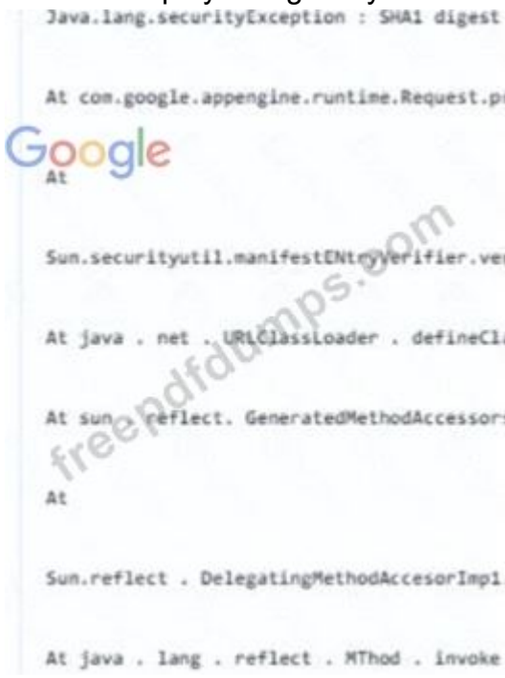

A. Recompile the CLoakedServlet class using and MD5 hash instead of SHA1
B. Digitally sign all of your JAR files and redeploy your application.
C. Upload missing JAR files and redeploy your application
Answer: B (LEAVE A REPLY)

NEW QUESTION: 108
You need to migrate Hadoop jobs for your company's Data Science team without modifying the underlying infrastructure. You want to minimize costs and infrastructure management effort. What should you do?
A. Create a Dataproc cluster using standard worker instances.
B. Create a Dataproc cluster using preemptible worker instances.
C. Manually deploy a Hadoop cluster on Compute Engine using standard instances.
D. Manually deploy a Hadoop cluster on Compute Engine using preemptible instances.
Answer: (SHOW ANSWER)
Reference: https://cloud.google.com/architecture/hadoop/hadoop-gcp-migration-jobs

NEW QUESTION: 109
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: A (LEAVE A REPLY)
Explanation
https://cloud.google.com/blog/products/storage-data-transfer/uploading-images-directly-to-cloud-storage-by-usin

NEW QUESTION: 110
Your company creates rendering software which users can download from the company website. Your company has customers all over the world. You want to minimize latency for all your customers. You want to follow Google-recommended practices.
How should you store the files?
A. Save the files in a Multi-Regional Cloud Storage bucket.
B. Save the files in a Regional Cloud Storage bucket, one bucket per zone of the region.
C. Save the files in multiple Regional Cloud Storage buckets, one bucket per zone per region.
D. Save the files in multiple Multi-Regional Cloud Storage buckets, one bucket per multi-region.

Answer: A (LEAVE A REPLY)
Explanation
https://cloud.google.com/storage/docs/locations#location-mr

NEW QUESTION: 111
You set up an autoscaling instance group to serve web traffic for an upcoming launch. After configuring the instance group as a backend service to an HTTP(S) load balancer, you notice that virtual machine (VM) instances are being terminated and re-launched every minute. The instances do not have a public IP address.
You have verified the appropriate web response is coming from each instance using the curl command. You want to ensure the backend is configured correctly. What should you do?
A. Ensure that a firewall rule exists to allow source traffic on HTTP/HTTPS to reach the load balancer.
B. Assign a public IP to each instance and configure a firewall rule to allow the load balancer to reach the instance public IP.
C. Ensure that a firewall rule exists to allow load balancer health checks to reach the instances in the instance group.
D. Create a tag on each instance with the name of the load balancer. Configure a firewall rule with the name of the load balancer as the source and the instance tag as the destination.
Answer: (SHOW ANSWER)
Explanation
https://cloud.google.com/vpc/docs/using-firewalls
The best practice when configuration a health check is to check health and serve traffic on the same port.
However, it is possible to perform health checks on one port, but serve traffic on another. If you do use two different ports, ensure that firewall rules and services running on instances are configured appropriately. If you run health checks and serve traffic on the same port, but decide to switch ports at some point, be sure to update both the backend service and the health check.
Backend services that do not have a valid global forwarding rule referencing it will not be health checked and will have no health status.
References: https://cloud.google.com/compute/docs/load-balancing/http/backend-service

NEW QUESTION: 112
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 Identity and Access Management (IAM) and Confidential Computing to trigger a Cloud Function.

B. Configure the deployment job to notify a Pub/Sub queue that triggers a Cloud Function.
C. Set up Cloud Tasks and a Cloud Storage bucket that triggers a Cloud Function.
D. Set up a Cloud Logging sink and a Cloud Storage bucket that triggers a Cloud Function.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 113
Your company has an application running on App Engine that allows users to upload music files and share them with other people. You want to allow users to upload files directly into Cloud Storage from their browser session. The payload should not be passed through the backend. What should you do?
A. 1. Assign the Cloud Storage WRITER role to users who upload files.
2. Use App Engine default credentials to sign requests against Cloud Storage.
B. 1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App Engine application is an allowed origin.
2. Use the Cloud Storage Signed URL feature to generate a POST URL.
C. 1. Use the Cloud Storage Signed URL feature to generate a POST URL.
2. Use App Engine default credentials to sign requests against Cloud Storage.
D. 1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App Engine application is an allowed origin.
2. Assign the Cloud Storage WRITER role to users who upload files.
Answer: (SHOW ANSWER)

NEW QUESTION: 114
For this question, refer to the Dress4Win case study.
Dress4Win has asked you to recommend machine types they should deploy their application servers to. How should you proceed?
A. Identify the number of virtual cores and RAM associated with the application server virtual machines align them to a custom machine type in the cloud, monitor performance, and scale the machine types up until the desired performance is reached.
B. Recommend that Dress4Win deploy into production with the smallest instances available, monitor them over time, and scale the machine type up until the desired performance is reached.
C. Recommend that Dress4Win deploy application servers to machine types that offer the highest RAM to CPU ratio available.
D. Perform a mapping of the on-premises physical hardware cores and RAM to the nearest machine types in the cloud.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 115
You need to deploy an application on Google Cloud that must run on a Debian Linux environment. The application requires extensive configuration in order to operate correctly. You want to ensure that you can install Debian distribution updates with minimal manual intervention whenever they become available. What should you do?

A. Create a Compute Engine instance template using the most recent Debian image. 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.
B. Create a Debian-based Compute Engine instance, install and configure the application, and use OS patch management to install available updates.
C. 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.
D. 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.
Answer: B (LEAVE A REPLY)
Reference: https://cloud.google.com/compute/docs/os-patch-management

NEW QUESTION: 116
For this question, refer to the TerramEarth case study.
You start to build a new application that uses a few Cloud Functions for the backend. One use case requires a Cloud Function func_display to invoke another Cloud Function func_query. You want func_query only to accept invocations from func_display. You also want to follow Google's recommended best practices. What should you do?
A. Create a token and pass it in as an environment variable to func_display. When invoking func_query, include the token in the request Pass the same token to func _query and reject the invocation if the tokens are different.
B. Make func_query 'Require authentication.' Create a unique service account and associate it to func_display. Grant the service account invoker role for func_query. Create an id token in func_display and include the token to the request when invoking func_query.
C. Make func _query 'Require authentication' and only accept internal traffic. Create those two functions in the same VPC. Create an ingress firewall rule for func_query to only allow traffic from func_display.
D. Create those two functions in the same project and VPC. Make func_query only accept internal traffic.Create an ingress firewall for func_query to only allow traffic from func_display. Also, make sure both functions use the same service account.
Answer: (SHOW ANSWER)
Explanation
https://cloud.google.com/functions/docs/securing/authenticating#authenticating_function_to_function_calls

NEW QUESTION: 117
For this question, refer to the TerramEarth case study.
TerramEarth has equipped unconnected trucks with servers and sensors to collet telemetry data. Next year they want to use the data to train machine learning models. They want to store this data in the cloud while reducing costs. What should they do?

A. Have the vehicle' computer compress the data in hourly snapshots, and store it in a Google Cloud storage (GCS) Nearline bucket.
B. Push the telemetry data in Real-time to a streaming dataflow job that compresses the data, and store it in Google BigQuery.
C. Push the telemetry data in real-time to a streaming dataflow job that compresses the data, and store it in Cloud Bigtable.
D. Have the vehicle's computer compress the data in hourly snapshots, a Store it in a GCS Coldline bucket.
Answer: D (LEAVE A REPLY)
Explanation
Coldline Storage is the best choice for data that you plan to access at most once a year, due to its slightly lower availability, 90-day minimum storage duration, costs for data access, and higher per-operation costs. For example:
Cold Data Storage - Infrequently accessed data, such as data stored for legal or regulatory reasons, can be stored at low cost as Coldline Storage, and be available when you need it.
Disaster recovery - In the event of a disaster recovery event, recovery time is key. Cloud Storage provides low latency access to data stored as Coldline Storage.
References: https://cloud.google.com/storage/docs/storage-classes
Topic 1, TerramEarth Case Study
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

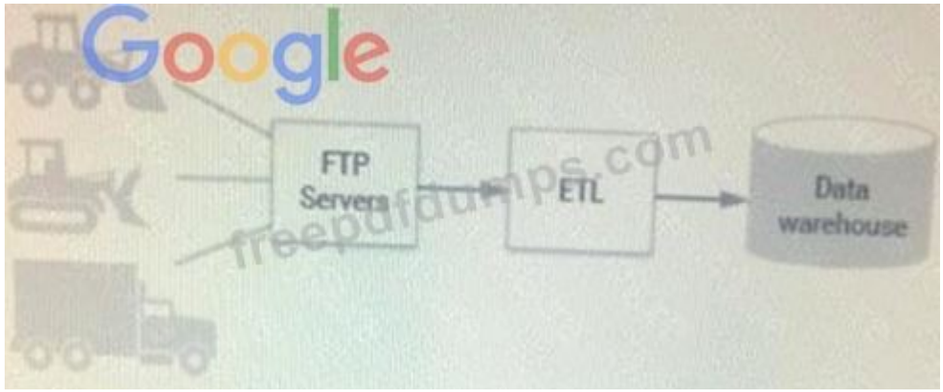

TerramEarth's existing architecture is composed of Linux-based systems that reside in a data center. These systems gzip CSV files from the field and upload via FTP, transform and aggregate them, and place the data in their data warehouse. Because this process takes time, aggregated reports are based 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 in the fast-growing agricultural business-to create compelling joint offerings for their customers 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 tower 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: 118

You are developing an application using different microservices that should remain internal to the cluster. You want to be able to configure each microservice with a specific number of replicas. You also want to

be able to address a specific microservice from any other microservice in a uniform way, regardless of the number of replicas the microservice scales to. You need to implement this solution on Google Kubernetes Engine. What should you do?

A. Deploy each microservice as a Deployment. Expose the Deployment in the cluster using a Service, and use the Service DNS name to address it from other microservices within the cluster.

B. Deploy each microservice as a Deployment. Expose the Deployment in the cluster using an Ingress, and use the Ingress IP address to address the Deployment from other microservices within the cluster.

C. Deploy each microservice as a Pod. Expose the Pod in the cluster using a Service, and use the Service DNS name to address the microservice from other microservices within the cluster.

D. Deploy each microservice as a Pod. Expose the Pod in the cluster using an Ingress, and use the Ingress IP address name to address the Pod from other microservices within the cluster.

Answer: A (LEAVE A REPLY)

Explanation

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

NEW QUESTION: 119

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

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

A. Google Kubernetes Engine with an SSL Ingress

B. Cloud IoT Core with public/private key pairs

C. Compute Engine with project-wide SSH keys

D. Compute Engine with specific SSH keys

Answer: A (LEAVE A REPLY)

Explanation

https://cloud.google.com/solutions/iot-overview

https://cloud.google.com/iot/quotas

NEW QUESTION: 120

TerramEarth has about 1 petabyte (PB) of vehicle testing data in a private data center. You want to move the data to Cloud Storage for your machine learning team. Currently, a 1-Gbps interconnect link is available for you. The machine learning team wants to start using the data in a month. What should you do?

A. Configure the Storage Transfer service from Google Cloud to send the data from your data center to Cloud Storage

B. Request Transfer Appliances from Google Cloud, export the data to appliances, and return the appliances to Google Cloud.

C. Export files to an encrypted USB device, send the device to Google Cloud, and request an import of the data to Cloud Storage

D. Make sure there are no other users consuming the 1 Gbps link, and use multi-thread transfer to upload the data to Cloud Storage.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 121
You are running a cluster on Kubernetes Engine to serve a web application. Users are reporting that a specific part of the application is not responding anymore. You notice that all pods of your deployment keep restarting after 2 seconds. The application writes logs to standard output. You want to inspect the logs to find the cause of the issue. Which approach can you take?
A. Review the Stackdriver logs for each Compute Engine instance that is serving as a node in the cluster.
B. Review the Stackdriver logs for the specific Kubernetes Engine container that is serving the unresponsive part of the application.
C. Review the Serial Port logs for each Compute Engine instance that is serving as a node in the cluster.
D. Connect to the cluster using gcloud credentials and connect to a container in one of the pods to read the logs.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 122
Your organization requires that metrics from all applications be retained for 5 years for future analysis in possible legal proceedings. Which approach should you use?
A. Grant the security team access to the logs in each Project.
B. Configure Stackdriver Monitoring for all Projects, and export to BigQuery.
C. Configure Stackdriver Monitoring for all Projects with the default retention policies.
D. Configure Stackdriver Monitoring for all Projects, and export to Google Cloud Storage.
Answer: D (LEAVE A REPLY)
Explanation
Overview of storage classes, price, and use cases https://cloud.google.com/storage/docs/storage-classes Why export logs? https://cloud.google.com/logging/docs/export/ StackDriver Quotas and Limits for Monitoring https://cloud.google.com/monitoring/quotas The BigQuery pricing. https://cloud.google.com/bigquery/pricing

NEW QUESTION: 123
Your company has a Google Workspace account and Google Cloud Organization Some developers in the company have created Google Cloud projects outside of the Google Cloud Organization You want to create an Organization structure that allows developers to create projects, but prevents them from modifying production projects You want to manage policies for all projects centrally and be able to set more restrictive policies for production projects You want to minimize disruption to users and developers when business needs change in the future You want to follow Google-recommended practices How should you design the Organization structure?
A. 1 Create a second Google Workspace account and Organization
2 Grant all developers the Project Creator IAM role on the new Organization
3 Move the developer projects into the new Organization
4 Set the policies for all projects on both Organizations.
5 Additionally set the production policies on the original Organization
B. 1 Create a folder under the Organization resource named "Production '
2 Grant all developers the Project Creator IAM role on the Organization 3. Move the developer projects into the Organization
4 Set the policies for all projects on the Organization
5 Additionally set the production policies on the 'Production" folder
C. 1 Create folders under the Organization resource named "Development" and Production'
2 Grant all developers the Project Creator IAM role on the ""Development1 folder 3. Move the developer projects into the "Development" folder
4 Set the policies for all projects on the Organization
5 Additionally set the production policies on the "Production" folder
D. 1 Designate the Organization for production projects only
2 Ensure that developers do not have the Project Creator IAM role on the Organization
3 Create development projects outside of the Organization using the developer Google Workspace accounts
4 Set the policies for all projects on the Organization
5 Additionally set the production policies on the individual production projects
Answer: (SHOW ANSWER)
Explanation
This option can help create an organization structure that allows developers to create projects, but prevents them from modifying production projects. Folders are containers for projects and other folders within Google Cloud organizations. Folders allow resources to be structured hierarchically and inherit policies from their parent resources. By creating folders under the organization resource named "Development" and
"Production", you can organize your projects by environment and apply different policies to them. By granting all developers the Project Creator IAM role on the "Development" folder, you can allow them to create projects under that folder, but not under the "Production" folder. By moving the developer projects into the

"Development" folder, you can ensure that they are subject to the policies set on that folder. By setting the policies for all projects on the organization, you can manage policies centrally and efficiently. By additionally setting the production policies on the "Production" folder, you can enforce more restrictive policies for production projects and prevent developers from modifying them. The other options are not optimal for this scenario, because they either create a second Google Workspace account and organization, which increases complexity and cost (A), or do not use folders to organize projects by environment, which makes it harder to manage policies and permissions (B, D). References:
https://cloud.google.com/resource-manager/docs/creating-managing-folders
https://cloud.google.com/architecture/framework/system-design

NEW QUESTION: 124
You are deploying an application on App Engine that needs to integrate with an on-premises database. For security purposes, your on-premises database must not be accessible through the public Internet. What should you do?
A. Deploy your application on App Engine standard environment and use App Engine firewall rules to limit access to the open on-premises database.
B. Deploy your application on App Engine standard environment and use Cloud VPN to limit access to the onpremises database.
C. Deploy your application on App Engine flexible environment and use App Engine firewall rules to limit access to the on-premises database.
D. Deploy your application on App Engine flexible environment and use Cloud VPN to limit access to the onpremises database.
Answer: (SHOW ANSWER)
Explanation
https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databases

NEW QUESTION: 125
You are developing a globally scaled frontend for a legacy streaming backend data API. This API expects events in strict chronological order with no repeat data for proper processing.
Which products should you deploy to ensure guaranteed-once FIFO (first-in, first-out) delivery of data?
A. Cloud Pub/Sub alone
B. Cloud Pub/Sub to Cloud DataFlow
C. Cloud Pub/Sub to Stackdriver
D. Cloud Pub/Sub to Cloud SQL
Answer: B (LEAVE A REPLY)
Explanation
Reference https://cloud.google.com/pubsub/docs/ordering

NEW QUESTION: 126

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

A. Create the Key object for each Entity and run a batch get operation

B. Create the Key object for each Entity and run multiple get operations, one operation for each entity

C. Use the identifiers to create a query filter and run a batch query operation

D. Use the identifiers to create a query filter and run multiple query operations, one operation for each entity

Answer: C (LEAVE A REPLY)

Explanation

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

NEW QUESTION: 127

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

A. Google Cloud Dataflow

B. Google Cloud Dataproc

C. Google Compute Engine

D. Google Container Engine

Answer: B (LEAVE A REPLY)

Explanation

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

References: https://cloud.google.com/dataproc/docs/resources/faq

NEW QUESTION: 128

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: C (LEAVE A REPLY)
Explanation
https://cloud.google.com/apis/design/versioning
All Google API interfaces must provide a major version number, which is encoded at the end of the protobuf package, and included as the first part of the URI path for REST APIs. If an API introduces a breaking change, such as removing or renaming a field, it must increment its API version number to ensure that existing user code does not suddenly break.

NEW QUESTION: 129
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 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
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 registered as a xinetd service in Linux and configure a Stackdnver endpoint check to call the service.
D. Create a shutdown script named k99.shutdown in the /etc/rc.6.d/ directory.
Answer: (SHOW ANSWER)

NEW QUESTION: 130
Your organization has a 3-tier web application deployed in the same network on Google Cloud Platform. Each tier (web, API, and database) scales independently of the others Network traffic should flow through the web to the API tier and then on to the database tier. Traffic should not flow between the web and the database tier.
How should you configure the network?
A. Add each tier to a different subnetwork.
B. Set up software based firewalls on individual VMs.
C. Add tags to each tier and set up routes to allow the desired traffic flow.
D. Add tags to each tier and set up firewall rules to allow the desired traffic flow.
Answer: (SHOW ANSWER)
Explanation
https://aws.amazon.com/blogs/aws/building-three-tier-architectures-with-security-groups/ Google Cloud Platform(GCP) enforces firewall rules through rules and tags. GCP rules and tags can be defined once and used across all regions.
References: https://cloud.google.com/docs/compare/openstack/
https://aws.amazon.com/it/blogs/aws/building-three-tier-architectures-with-security-groups/

NEW QUESTION: 131
Your company is running its application workloads on Compute Engine. The applications have been deployed in production, acceptance, and development environments. The production environment is business-critical and is used 24/7, while the acceptance and development environments are only critical during office hours.
Your CFO has asked you to optimize these environments to achieve cost savings during idle times. What should you do?
A. Create a shell script that uses the gcloud command to change the machine type of the development and acceptance instances to a smaller machine type outside of office hours. Schedule the shell script on one of the production instances to automate the task.
B. Use Cloud Scheduler to trigger a Cloud Function that will stop the development and acceptance environments after office hours and start them just before office hours.
C. Deploy the development and acceptance applications on a managed instance group and enable autoscaling.
D. Use regular Compute Engine instances for the production environment, and use preemptible VMs for the acceptance and development environments.
Answer: (SHOW ANSWER)
Reference: https://cloud.google.com/blog/products/it-ops/best-practices-for-optimizing-your-cloud-costs