Google.Professional-Cloud-Architect.v2026-11-03.q150

Exam Code:	Professional-Cloud-Architect
Exam Name:	Google Certified Professional - Cloud Architect (GCP)
Certification Provider:	Google
Free Question Number:	150
Version:	v2026-11-03
# of views:	114
# of Questions views:	1780
https://www.freecram.net/torrent/Google.Professional-Cloud-Architect.v2026-11-03.q150.html	

NEW QUESTION: 1
Refer to the Altostrat Media case study for the following solution regarding API management and cost control.
Altostrat is using Apigee for API management and wants to ensure their APIs are protected from overuse and abuse. You need to implement an Apigee feature to control the total number of API calls for cost management. What should you do?
A. Set up API key validation.
B. Integrate OAuth 2.0 authorization.
C. Configure Quota policies.
D. Activate XML threat protection.
Answer: (SHOW ANSWER)
According to the Apigee API Platform documentation , the Quota policy is the primary mechanism used to enforce business-level constraints by limiting the number of request messages an API proxy allows over a specific period of time (such as a minute, hour, day, week, or month). For a media company like Altostrat, which aims to " unlock new revenue streams through targeted marketing " and " drive revenue growth, " managing API consumption is critical for cost predictability and monetization strategies.
Google Cloud distinguishes between Spike Arrest and Quota policies . While Spike Arrest is designed for operational safety-protecting backend systems from sudden traffic bursts-the Quota policy is specifically used to enforce business contracts, Service Level Agreements (SLAs), and cost management . By configuring a Quota policy, Altostrat can ensure that no single consumer or application exceeds the total volume of calls that would result in unexpected infrastructure costs or exceed the limits of their current subscription tier.
Options A and B (API Key and OAuth) are essential for authentication and authorization -identifying who is making the call-but they do not inherently limit the volume of calls. Option D (XML threat protection) is a security measure to prevent malformed payload attacks. Therefore, to meet the specific business requirement of " Reliability and cost management, " implementing Quota policies within Apigee is the architecturally correct choice to prevent API abuse and manage operational expenditure.

NEW QUESTION: 2
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: (SHOW ANSWER)
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: 3
You are designing the architecture for a stateless, public-facing, containerized web application. The marketing department uses this web application to run unpredictable flash sales that cause sudden, massive spikes in user traffic. The application is a critical revenue driver and must always be available and responsive to users. Your design must automatically handle high traffic fluctuations and allow for new application versions to be deployed with zero downtime and a safe, fast rollback capability. What should you do?
A. Deploy the application to a regional Managed Instance Group (MIG). Configure a CPU-based autoscaling policy and use the MIG ' s rolling update feature for deployments.
B. Deploy the application as a Cloud Run service. Configure request-based auto-scaling and use gradual rollouts by splitting traffic between revisions.
C. Deploy the application to a single, large Compute Engine instance and use a startup script to pull and run the latest container image on boot.
D. Deploy the application to a Google Kubernetes Engine (GKE) cluster. Use the Recreate deployment strategy to roll out new application versions.
Answer: (SHOW ANSWER)
For stateless, containerized applications facing " sudden, massive spikes, " Cloud Run is the preferred serverless solution. According to Cloud Run documentation , it offers superior scaling speed compared to VMs or GKE, as it can scale from zero to thousands of instances in seconds based on request-based metrics (concurrency) rather than just CPU utilization.

Cloud Run natively supports Revisions , allowing for zero-downtime deployments . You can perform " Blue-Green " or " Canary " deployments by splitting traffic between a new revision and an old one. If a new version fails, the rollback is nearly instantaneous (switching traffic back to the previous revision). Option A (MIGs) scales significantly slower than Cloud Run. Option D (GKE Recreate) involves downtime during deployments, and Option C is a single point of failure that cannot handle traffic spikes.

NEW QUESTION: 4
The operations team in your company wants to save Cloud VPN log events (or one year You need to configure the cloud infrastructure to save the logs What should you do?
A. Set up a filter in Cloud Logging and a Cloud Storage bucket as an export target for the logs you want to save
B. Enable the Compute Engine API and then enable logging on the firewall rules that match the traffic you want to save
C. Set up a Cloud Logging Dashboard titled Cloud VPN Logs, and then add a chart that queries for the VPN metrics over a one-year time period
D. Set up a filter in Cloud Logging and a topic in Pub/Sub to publish the logs
Answer: (SHOW ANSWER)

NEW QUESTION: 5
A retail company ' s most critical application is its online payment processing system. The business has a requirement that the system must be able to survive a complete zonal outage while minimizing cost. You need a design solution that can handle a zonal failure. What should you do?
A. Deploy the application in an active-active configuration using managed instance groups (MIGs) in two different regions, fronted by a global external HTTP(S) Load Balancer and backed by a multi-regional database like Spanner.
B. Deploy the application on a regional MIG to provide high availability across multiple zones in the primary region.
C. Configure the regional MIG to use only Spot VMs to aggressively minimize operational costs while maintaining high availability.
D. Deploy the application on Compute Engine instances across multiple regions, and rely on daily snapshots for recovery to achieve the lowest possible cost.
Answer: (SHOW ANSWER)
The Google Cloud Architecture Framework: Reliability section emphasizes that for applications requiring protection against zonal failures , the most cost-effective solution is a Regional Managed Instance Group (MIG) . A Regional MIG automatically distributes VM instances across multiple zones within a single region. If one zone fails, the MIG's auto-healing and cross-zone load balancing ensure that instances in the remaining healthy zones continue to serve traffic.
Option B is the correct choice because it satisfies the " survive a zonal outage " requirement while " minimizing cost. " Multi-region deployments (Option A) provide higher availability (surviving regional outages) but significantly increase costs due to inter-region data transfer and the expense of multi-region database configurations like Spanner. Spot VMs (Option C) are unsuitable for a " critical payment

processing system " because they can be preempted at any time, risking transaction failures. Snapshot recovery (Option D) results in high Recovery Time Objectives (RTO), failing the performance expectations of a critical retail application. By using a Regional MIG with a regional database (like Cloud SQL with High Availability enabled), the company achieves robust zonal resilience at a much lower price point than a global multi-region architecture.

NEW QUESTION: 6
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: (SHOW ANSWER)
https://cloud.google.com/service-mesh/docs/overview
https://cloud.google.com/service-mesh/docs/observability/slo-overview

NEW QUESTION: 7
You deploy your custom Java application to Google App Engine. It fails to deploy and gives you the following stack trace.
What should you do?
A. Recompile the CLoakedServlet class using and MD5 hash instead of SHA1
B. Upload missing JAR files and redeploy your application.
C. Digitally sign all of your JAR files and redeploy your application
Answer: (SHOW ANSWER)

NEW QUESTION: 8
You are a cloud architect for a financial services company that is subject to strict regulatory compliance for data residency and business continuity. Your primary transaction processing application runs on a Compute Engine instance in europe-west3, and its critical data is stored on the attached regional Persistent Disk. To meet business continuity requirements, you must ensure the application can be recovered in a different region with minimal data loss in the event of a regional outage and follow Google-recommended best practices.
What should you do?
A. Create a snapshot schedule for the Persistent Disk that copies the snapshots to a different region.

B. Use Backup and DR Service to create backups, and store them in a different zone within the same source region.
C. Configure the regional Persistent Disk for synchronous replication.
D. Create a custom image from the instance, and configure it to be stored in a multi-regional location.
Answer: (SHOW ANSWER)
Regional Persistent Disks provide synchronous data replication across two distinct zones within a single Google Cloud region. While this mitigates local zonal failures, it provides zero protection against a full re gional outage. To ensure cross-region business continuity and comply with data residency boundaries, backup data must be stored externally in a distinct target region.
Google-recommended best practice for cross-region block storage recovery involves implementing automated Snapshot Schedules. Snapshots are incremental by nature, capturing disk deltas to reduce overhead, and they can be configured to replicate to a specific, custom secondary region. This satisfies data residency requirements by keeping data confined to pre-approved regulatory boundaries while decoupling the data from the source region ' s infrastructure. In a disaster scenario, these snapshots can quickly provision new Persistent Disks in the target region. Custom images only capture boot configurations rather than dynamic application state data, and local zone backups fail the regional resiliency test.
References: Cloud Architect (GCP) Study Guide topics: " Storage Resiliency " , " Persistent Disk Snapshots Architecture " , " Cross-Region Disaster Recovery " , and " Data Residency Compliance " .

NEW QUESTION: 9
Your company plans to migrate a multi-petabyte data set to the cloud. The data set must be available 24hrs a day. Your business analysts have experience only with using a SQL interface. How should you store the data to optimize it for ease of analysis?
A. Load data into Google BigQuery.
B. Insert data into Google Cloud SQL.
C. Put flat files into Google Cloud Storage.
D. Stream data into Google Cloud Datastore.
Answer: (SHOW ANSWER)
Google Big Query is for multi peta byte storage , HA(High availability) which means 24 hours, SQL interface .
https://medium.com/google-cloud/the-12-components-of-google-bigquery-c2b49829a7c7 h
ttps ://cloud.google.com/solutions/bigquery-data-warehouse
https://cloud.google.com/bigquery/
BigQuery is Google ' s serverless, highly scalable, low cost enterprise data warehouse designed to make all your data analysts productive. Because there is no infrastructure to manage, you can focus on analyzing data to find meaningful insights using familiar SQL and you don ' t need a database administrator.
BigQuery enables you to analyze all your data by creating a logical data warehouse over managed, columnar storage as well as data from object storage, and spreadsheets.
References: https: //cloud.google.com/bigquery/

NEW QUESTION: 10

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

Freecram.net

A. Insert logs into Google Cloud Bigtable.
B. Load logs into Google Cloud SQL.
C. Upload log files into Google Cloud Storage.
D. Import logs into Google Stackdriver.
E. Load logs into Google BigQuery.
Answer: (SHOW ANSWER)

NEW QUESTION: 11
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 PII 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 PII 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 PII 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 PII 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 PII data
Answer: (SHOW ANSWER)
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: 12
For this question, refer to the Dress4Win case study. Dress4Win is expected to grow to 10 times its size in 1 year with a corresponding growth in data and traffic that mirrors the existing patterns of usage. The CIO has set the target of migrating production infrastructure to the cloud within the next 6 months. How will you configure the solution to scale for this growth without making major application changes and still maximize the ROI?
A. Migrate the web application layer to App Engine, and MySQL to Cloud Datastore, and NAS to Cloud Storage. Deploy RabbitMQ, and deploy Hadoop servers using Deployment Manager.
B. Migrate RabbitMQ to Cloud Pub/Sub, Hadoop to BigQuery, and NAS to Compute Engine with Persistent Disk storage. Deploy Tomcat, and deploy Nginx using Deployment Manager.
C. Implement managed instance groups for Tomcat and Nginx. Migrate MySQL to Cloud SQL, RabbitMQ to Cloud Pub/Sub, Hadoop to Cloud Dataproc, and NAS to Compute Engine with Persistent Disk storage.

D. Implement managed instance groups for the Tomcat and Nginx. Migrate MySQL to Cloud SQL, RabbitMQ to Cloud Pub/Sub, Hadoop to Cloud Dataproc, and NAS to Cloud Storage.

Answer: (SHOW ANSWER)

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: 13
Your company has successfully migrated to the cloud and wants to analyze their data stream to optimize operations. They do not have any existing code for this analysis, so they are exploring all their options. These options include a mix of batch and stream processing, as they are running some hourly jobs and live-processing some data as it comes in. Which technology should they use for this?
A. Google Cloud Dataproc
B. Google Cloud Dataflow
C. Google Container Engine with Bigtable
D. Google Compute Engine with Google BigQuery
Answer: B (LEAVE A REPLY)
Dataflow is for processing both the Batch and Stream.
Cloud Dataflow is a fully-managed service for transforming and enriching data in stream (real time) and batch (historical) modes with equal reliability and expressiveness -- no more complex workarounds or compromises needed.

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

NEW QUESTION: 14
A healthcare organization is migrating its patient data analytics platform to Google Cloud. The system runs on a large number of VMs and handles highly sensitive patient data. To comply with regulations, the data must be encrypted while in use. Additionally, a primary security requirement is to prevent data exfiltration by ensuring the platform ' s managed services can only be accessed from within a secure network boundary and are blocked from accessing other Google Cloud APIs. The organization has a short timeline for the migration and wants to minimize downtime and is using Migrate to Virtual Machines for the migration. What should you do?
A. Use Confidential VMs for the migrated workloads. Use Identity-Aware Proxy (IAP) to secure all SSH and RDP access to the VMs.
B. Use Shielded VMs for the migrated workloads. Configure firewall rules to restrict access to the backend services. Configure a VPC Service Controls perimeter to protect the backend services.
C. Use Migrate to Virtual Machines to create a runbook automation for a wave-based migration. Apply customer-managed encryption keys (CMEK) to the persistent disks of the target VMs. Configure a VPC Service Controls perimeter to protect the backend services.
D. Use Confidential VMs for the migrated workloads. Configure a VPC Service Controls perimeter to protect the backend services.
Answer: (SHOW ANSWER)
Regulatory compliance for encrypting data while in use necessitates the deployment of Confidential VMs. Confidential Computing leverages specialized hardware extensions (such as AMD SEV) to keep memory contents encrypted during CPU execution, protecting data from runtime inspection. Standard or Shielded VMs only encrypt data at rest and in transit, falling short of the in-use cryptographic security requirements. To address the data exfiltration risk and satisfy the secure network boundary requirement, a VPC Service Controls (VPC-SC) security perimeter must be wrapped around the Google Cloud managed services. VPC Service Controls establishes a logical perimeter that restricts access to supported API endpoints to requests originating from within specified VPC networks, effectively neutralizing data exfiltration attempts through unauthorized external APIs. Combining Confidential VMs for runtime memory encryption with VPC Service Controls for strict API perimeter containment directly satisfies both regulatory and technical goals.
References: Cloud Architect (GCP) Study Guide topics: " Confidential Computing Architecture " , " VPC Service Controls (VPC-SC) Perimeters " , " Data Exfiltration Prevention " , and " Data in Use Encryption " .

NEW QUESTION: 15
A production database virtual machine on Google Compute Engine has an ext4-formatted persistent disk for data files The database is about to run out of storage space How can you remediate the problem with the least amount of downtime?
A. In the Cloud Platform Console, increase the size of the persistent disk and use the resize2fs command in Linux.
B. Shut down the virtual machine, use the Cloud Platform Console to increase the persistent disk size, then restart the virtual machine.

C. In the Cloud Platform Console, increase the size of the persistent disk and verify the new space is ready to use with the fdisk command in Linux.
D. In the Cloud Platform Console, create a new persistent disk attached to the virtual machine, format and mount it, and configure the database service to move the files to the new disk.
E. In the Cloud Platform Console, create a snapshot of the persistent disk, restore the snapshot to a new larger disk, unmount the old disk, mount the new disk, and restart the database service.
Answer: (SHOW ANSWER)
On Linux instances, connect to your instance and manually resize your partitions and file systems to use the additional disk space that you added.
Extend the file system on the disk or the partition to use the added space. If you grew a partition on your disk, specify the partition. If your disk does not have a partition table, specify only the disk ID.
sudo resize2fs /dev/[DISK_ID][PARTITION_NUMBER]
where [DISK_ID] is the device name and [PARTITION_NUMBER] is the partition number for the device where you are resizing the file system.
References: https://cloud.google.com/compute/docs/disks/add-persistent-disk

NEW QUESTION: 16
You are designing a central, automated infrastructure deployment process tor your organization using Terraform and Cloud Build The security team prohibits the use of long-lived, static service account keys in any CI/CD pipeline Additionally, while developers can propose infrastructure changes for peer review, they must not have permissions to directly apply changes in the production project. You need to design a secure and automated workflow for applying Terraform changes that meets the security team ' s requirements and ensures proper governance. What should you do?
A. Configure the Cloud Build pipeline to use service account impersonation. Set up a trigger that automatically runs terraform apply when a pull request is merged.
B. Create a privileged service account and store its JSON key in Secret Manager Configure the Cloud Build pipeline to fetch this key during execution to authenticate Terraform.
C. Configure the pipeline to only run terraform plan. After a pull request Is approved, have an authorized developer run terraform apply from a secured workstation
D. Use service account impersonation in Cloud Build. Configure the pipeline to run terraform plan on pull requests, and require manual approval before running terraform apply.
Answer: (SHOW ANSWER)
Security best practices in Google Cloud dictate that you should avoid static JSON keys at all costs.
* Impersonation: By using Service Account Impersonation (Option D), the Cloud Build service account " acts as " a more privileged Terraform service account only during the build. This uses short- lived tokens that expire automatically.
* Governance: Running terraform plan on a pull request allows peers to review the exact changes before they happen. Requiring a manual approval step for terraform apply ensures that a human gatekeeper validates the change before it impacts the production environment.

Valid Professional-Cloud-Architect Dumps shared by ExamDiscuss.com for Helping Passing Professional-Cloud-Architect Exam! ExamDiscuss.com now offer the **newest Professional-Cloud-Architect exam dumps**, the ExamDiscuss.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Professional-Cloud-Architect dumps with Test Engine here: https://www.examdiscuss.com/Google/exam/Professional-Cloud-Architect/premium/ (**353** Q&As Dumps, **35%OFF** **Special Discount Code: freecram**)

NEW QUESTION: 17
Your company uses Google Kubernetes Engine (GKE) as a platform for all workloads. Your company has a single large GKE cluster that contains batch, stateful, and stateless workloads. The GKE cluster is configured with a single node pool with 200 nodes. Your company needs to reduce the cost of this cluster but does not want to compromise availability. What should you do?
A. Create a second GKE cluster for the batch workloads only. Allocate the 200 original nodes across both clusters.
B. Configure a HorizontalPodAutoscaler for all stateless workloads and for all compatible stateful workloads. Configure the cluster to use node auto scaling.
C. Configure CPU and memory limits on the namespaces in the cluster. Configure all Pods to have a CPU and memory limits.
D. Change the node pool to use spot VMs.
Answer: (SHOW ANSWER)
One way to reduce the cost of a Google Kubernetes Engine (GKE) cluster without compromising availability is to use horizontal pod autoscalers (HPA) and node auto scaling. HPA allows you to automatically scale the number of Pods in a deployment based on the resource usage of the Pods. By configuring HPA for stateless workloads and for compatible stateful workloads, you can ensure that the number of Pods is automatically adjusted based on the actual resource usage, which can help to reduce costs. Node auto scaling allows you to automatically add or remove nodes from the node pool based on the resource usage of the cluster. By configuring node auto scaling, you can ensure that the cluster has the minimum number of nodes needed to meet the resource requirements of the workloads, which can also help to reduce costs.

NEW QUESTION: 18
You are a cloud architect for a global financial services company. Your primary production environment is in a VPC network in Google Cloud. Your company also operates a secondary VPC network in Google Cloud and a separate, smaller analytics workload in another major cloud provider. You need to design a network architecture that provides reliable, high-bandwidth connectivity between all three environments. Your goal is to use Google ' s private backbone as much as possible for all intercloud and inter-VPC traffic to ensure low latency and simplify network management. What should you do?
A. Configure Cross-Cloud Interconnect to connect your primary VPC network to the other cloud provider ' s VPC. Configure Network Connectivity Center in your primary Google Cloud project. Create spokes for the Cross-Cloud Interconnect connection and the secondary VPC network.

B. Establish a Dedicated Interconnect connection with a 99.99% SLA to the other cloud provider. Use VPC Network Peering to connect the primary and secondary VPC networks. Configure BGP sessions to exchange routes.
C. Create dedicated HA VPN tunnels from your primary VPC network to both the secondary VPC network and the other cloud provider ' s VPC. Configure dynamic routing with Cloud Router to manage traffic between the environments.
D. Configure VPC Network Peering between the primary and secondary VPC networks. Establish a Cross-Cloud Interconnect connection from the primary VPC network to the other cloud provider ' s VPC.Configure routes to direct traffic from the secondary VPC network through the primary VPC to the other cloud.
Answer: (SHOW ANSWER)
Comprehensive and Detailed 150 to 250 words of Explanation From [Cloud Architect (GCP) Course Guide /topics]: Cross-Cloud Interconnect provides high-bandwidth, private, and SLA-backed physical connectivity between Google Cloud and other hyperscale cloud service providers (such as AWS and Azure) without requiring intermediary colocation facilities. This leverages Google ' s global private fiber infrastructure directly to minimize latency and improve reliability for multicloud deployments.
Connecting multiple networks-specifically integrating a secondary Google Cloud VPC and an external cloud provider into a unified, transit-capable topology-requires Network Connectivity Center (NCC). Standard VPC Network Peering is strictly non-transitive; traffic from a peered secondary VPC cannot traverse the primary VPC to exit out of an Interconnect attachment. Network Connectivity Center addresses this constraint by acting as a centralized hub-and-spoke orchestration model over Google ' s backbone. By registering the Cross-Cloud Interconnect VLAN attachments and the secondary VPC network as spokes attached to an NCC hub, Google Cloud facilitates full, managed mesh routing and transit connectivity between the environments, centralizing management and keeping traffic on Google ' s private network.
References: Cloud Architect (GCP) Study Guide topics: " Cross-Cloud Interconnect Architecture " , " Network Connectivity Center (NCC) Hub and Spoke Topologies " , and " Transitive Routing Constraints in VPC Peering " .

NEW QUESTION: 19
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: 20

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

Freecram.net

A. Create an Organizational Policy with a constraint to allow external IP addresses only on the frontend Compute Engine instances.

B. Revoke the compute.networkAdmin role from all users in the project with front end instances.

C. Create an Identity and Access Management (IAM) policy that maps the IT staff to the compute. networkAdmin role for the organization.

D. Create a custom Identity and Access Management (IAM) role named GCE_FRONTEND with the compute.addresses.create permission.

Answer: (SHOW ANSWER)

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

NEW QUESTION: 21

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

How should you set up the connection?

A. Create a VPC and connect it to your on-premises data center using Dedicated Interconnect.

B. Create a VPC and connect it to your on-premises data center using a single Cloud VPN.

C. Create a Cloud Content Delivery Network (Cloud CDN) and connect it to your on-premises data center using Dedicated Interconnect.

D. Create a Cloud Content Delivery Network (Cloud CDN) and connect it to your on-premises datacenter using a single Cloud VPN.

Answer: (SHOW ANSWER)

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

NEW QUESTION: 22

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: (SHOW ANSWER)

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

NEW QUESTION: 23
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 Serial Port logs for each Compute Engine instance that is serving as a node in the cluster.
B. Connect to the cluster using gcloud credentials and connect to a container in one of the pods to read the logs.
C. Review the Stackdriver logs for the specific Kubernetes Engine container that is serving the unresponsive part of the application.
D. Review the Stackdriver logs for each Compute Engine instance that is serving as a node in the cluster.
Answer: (SHOW ANSWER)

NEW QUESTION: 24
For this question, refer to the TerramEarth case study.
To speed up data retrieval, more vehicles will be upgraded to cellular connections and be able to transmit data to the ETL process. The current FTP process is error-prone and restarts the data transfer from the start of the file when connections fail, which happens often. You want to improve the reliability of the solution and minimize data transfer time on the cellular connections. What should you do?
A. Use one Google Container Engine cluster of FTP servers. Save the data to a Multi-Regional bucket. Run the ETL process using data in the bucket.
B. Use multiple Google Container Engine clusters running FTP servers located in different regions. Save the data to Multi-Regional buckets in us, eu, and asia. Run the ETL process using the data in the bucket.
C. Directly transfer the files to different Google Cloud Multi-Regional Storage bucket locations in us, eu, and asia using Google APIs over HTTP(S). Run the ETL process using the data in the bucket.
D. Directly transfer the files to a different Google Cloud Regional Storage bucket location in us, eu, and asia using Google APIs over HTTP(S). Run the ETL process to retrieve the data from each Regional bucket.
Answer: (SHOW ANSWER)
https://cloud.google.com/storage/docs/locations

NEW QUESTION: 25
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: (SHOW ANSWER)

https://clo ud.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: 26

Your ecommerce platform uses a regional Cloud SQL for PostgreSQL database to store critical order information. The business requests a recovery time objective (RTO) of less than 10 minutes and a recovery point objective (RPO) of 2 minutes to ensure business continuity in the event of a full regional outage. You need to design a disaster recovery strategy for the Cloud SQL database that meets the business ' s strict RTO and RPO requirements. Your design must avoid operational complexity. What should you do?

A. Configure the Cloud SQL instance with a cross-region read replica. In a disaster, promote the read replica to a standalone, primary instance.

B. Use Database Migration Service to continuously replicate the database to another instance in a different region. In a disaster, redirect application traffic to the replica.

C. Schedule hourly automated backups of the Cloud SQL instance to a multi-regional Cloud Storage bucket. In a disaster, restore the latest backup to a new instance in a different region.

D. Configure the primary instance for high availability (HA). In the event of a regional outage, trigger a failover to the standby instance.

Answer: (SHOW ANSWER)

A full regional outage disables an entire Google Cloud region, including all internal availability zones. Standard High Availability (HA) configuration for Cloud SQL operates exclusively as an intra-region, cross-zone failover mechanism; it cannot survive a comprehensive regional disaster. Similarly, point-in-time recovery utilizing automated hourly backups cannot satisfy a strict 2-minute RPO requirement, as data generated between the last backup and the outage would be lost.

To satisfy an RPO of 2 minutes and an RTO of under 10 minutes during a total regional loss, a cross-region read replica must be maintained. Cloud SQL uses asynchronous replication to update cross-region replicas, which continuously mirrors transactions with minimal lag, safely preserving data in a geographically separated region. When the primary region drops, administrative promotion of the cross-region read replica into a standalone read-write primary instance can be completed in minutes, successfully meeting both the availability goals and minimizing complexity.

References: Cloud Architect (GCP) Study Guide topics: " Cloud SQL High Availability vs Disaster Recovery " , " Cross-Region Replication Mechanics " , and " Promoting Read Replicas for Business Continuity " .

NEW QUESTION: 27

Your designing the access control strategy for a new system that will store highly sensitive financial reports in a series of Cloud Storage buckets. The design must allow members of Google Groups to access only the

reports they are authorized to see. All access must be directly tied to an employee ' s corporate identity for clear audit trails and the access model must be simple to operate and scalable. What should you do?
A. Generate and distribute time-limited signed URLs to grant users temporary, per-object access.
B. Grant the IAM role roles/storage.objectAdmin to all authorized Google Groups to simplify permission management.
C. Enable uniform bucket-level access on the buckets, and grant predefined Storage IAM roles to the Google Groups.
D. Maintain fine-grained access control, and use default object ACLs to grant the required permissions to the Google Groups.
Answer: (SHOW ANSWER)
According to Google Cloud Storage Security documentation, Uniform Bucket-level Access is the industry-standard recommendation for securing sensitive enterprise data. When uniform access is enabled, it disables Access Control Lists (ACLs) entirely, forcing all permissions to be managed via Cloud IAM .
Using IAM roles (such as roles/storage.objectViewer) granted to Google Groups provides several advantages:
* Auditability: Every access request is logged with the user ' s corporate identity in Cloud Audit Logs , meeting the compliance requirement.
* Simplicity & Scalability: Permissions are managed in one place (the bucket level) rather than on individual files. As employees join or leave the company, their access is automatically updated by modifying their membership in the Google Group.
* Security: It prevents " permissions creep " and accidental exposure that often occurs with fine-grained ACLs (Option D).
Option A (Signed URLs) is over-engineered for internal employee access and harder to audit. Option B is a violation of the Principle of Least Privilege , as objectAdmin allows users to delete and overwrite files, which is inappropriate for financial report viewing.

NEW QUESTION: 28
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: (SHOW ANSWER)
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: 29
Your web application uses Google Kubernetes Engine to manage several workloads. One workload requires a consistent set of hostnames even after pod scaling and relaunches.
Which feature of Kubernetes should you use to accomplish this?
A. StatefulSets
B. Role-based access control
C. Container environment variables
D. Persistent Volumes
Answer: (SHOW ANSWER)
https://kubernetes.io/docs/tutorials/stateful-application/basic-stateful-set/

NEW QUESTION: 30
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?
Freecram.net
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)
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-5c847ac8f9f2
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: 31
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)
In this scenario, a large organization has a central team that manages security and networking controls for the entire organization. Developers do not have permissions to make changes to any network or security settings defined by the security and networking team but they are granted permission to create resources such as virtual machines in shared subnets. To facilitate this the organization makes use of a shared VPC (Virtual Private Cloud). A shared VPC allows creation of a VPC network of RFC 1918 IP spaces that associated projects (service projects) can then use. Developers using the associated projects can create

VM instances in the shared VPC network spaces. The organization ' s network and security admins can create subnets, VPNs, and firewall rules usable by all the projects in the VPC network. https://cloud.google.com/iam/docs/job-functions/networking#single_team_manages_security_network_for_organization Reference: https://cloud.google.com/vpc/docs/shared-vpc

Valid Professional-Cloud-Architect Dumps shared by ExamDiscuss.com for Helping Passing Professional-Cloud-Architect Exam! ExamDiscuss.com now offer the **newest Professional-Cloud-Architect exam dumps**, the ExamDiscuss.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Professional-Cloud-Architect dumps with Test Engine here: https://www.examdiscuss.com/Google/exam/Professional-Cloud-Architect/premium/ (**353** Q&As Dumps, **35%OFF** **Special Discount Code:** **freecram**)

NEW QUESTION: 32
You are deploying an application to Google Cloud. The application is part of a system. The application in Google Cloud must communicate over a private network with applications in a non-Google Cloud environment. The expected average throughput is 200 kbps. The business requires:
* 99.99% system availability
* cost optimization
You need to design the connectivity between the locations to meet the business requirements. What should you provision?
A. A Classic Cloud VPN gateway connected with one tunnel to an on-premises VPN gateway.
B. A Classic Cloud VPN gateway connected with two tunnels to an on-premises VPN gateway.
C. An HA Cloud VPN gateway connected with two tunnels to an on-premises VPN gateway.
D. Two HA Cloud VPN gateways connected to two on-premises VPN gateways. Configure each HA Cloud VPN gateway to have two tunnels, each connected to different on-premises VPN gateways.
Answer: (SHOW ANSWER)
https://cloud.google.com/network-connectivity/docs/vpn/concepts /topologies#configurations_that_support_9999_availability

NEW QUESTION: 33
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 Cloud SQL. Restore the latest backup. Use kubect1 to restart all pods.
B. Validate that the Service Account used by the Cloud SQL proxy container still has the Cloud Build Editor role.

C. In the GCP Console, navigate to Stackdriver Logging. Consult logs for Kubernetes Engine and Cloud SQL.
D. Use gcloud sql instances restart.
Answer: (SHOW ANSWER)

NEW QUESTION: 34
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)
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: 35
You are migrating a large, on-premises application to Google Cloud. The application consists of several interconnected virtual machines. You want to create a detailed migration plan to ensure a smooth migration with minimal effort. You need to understand the existing environment, identify dependencies, and estimate the total cost of ownership (TCO) in the cloud. What should you do?
A. Use Config Connector to declare the desired state of your Google Cloud resources in Kubernetes-style manifests.
B. Use the Google Cloud Migration Center to perform an automated discovery and assessment of the on-premises environment.
C. Use the Google Cloud pricing calculator to input the specifications of your on-premises servers and receive a TCO estimate.
D. Write a custom script to query the vSphere API for virtual machine information and then import it into BigQuery for analysis.
Answer: (SHOW ANSWER)
The Google Cloud Migration Center is the unified platform for end-to-end migration. According to the Cloud Migration Framework , the " Assess " phase is the most critical step for a successful migration.
The Migration Center provides automated discovery tools (like the mDiscovery tool) that scan your on-premises environment (vSphere, physical servers, etc.) to collect hardware specs, performance data, and

network dependencies . It then generates a TCO (Total Cost of Ownership) report comparing your current costs with optimized Google Cloud sizing. This replaces manual processes like Option D and Option C (which requires you to already know the specs). Option A (Config Connector) is used for managing existing cloud resources, not for assessing on-premises environments for migration.

NEW QUESTION: 36
You need to evaluate your team readiness for a new GCP project. You must perform the evaluation and create a skills gap plan incorporates the business goal of cost optimization. Your team has deployed two GCP projects successfully to date. What should you do?
A. Allocate budget for team training. Set a deadline for the new GCP project.
B. Allocate budget for team training. Create a roadmap for your team to achieve Google Cloud certification based on job role.
C. Allocate budget to hire skilled external consultants. Set a deadline for the new GCP project.
D. Allocate budget to hire skilled external consultants. Create a roadmap for your team to achieve Google Cloud certification based on job role.
Answer: B (LEAVE A REPLY)
https://services.google.com/fh/files/misc/cloud_center_of_excellence.pdf

NEW QUESTION: 37
You need to build a continuous delivery pipeline tor a containerized application in Google Cloud. You want to run all your tests in the pipeline to Improve your application ' s quality. What should you do?
A. 1. After the developers push the code to a central repository, trigger Cloud Build to run unit tests If all unit tests are successful, build the application container and push it to a central registry.
2. Trigger Cloud Build to deploy the container to a testing environment and run integration tests and acceptance tests
3. If all tests are successful, deploy the application to the production environment and run the smoke tests
B. 1. Automatically run unit tests in a local environment once the code is changed It all tests are successful, build a container.
2. Trigger Cloud Build to deploy the container to a testing environment, and run integration tests and acceptance tests.
3. If all tests are successful, tag the code as production ready. Trigger Cloud Build to build and deploy the container to the production environment.
C. 1. Run unit tests in the developer ' s local environment before committing and pushing the code to a central repository
2. After the code is pushed, trigger Cloud Build to build the application container and deploy the container to a testing environment, and run integration tests
3. If the integration tests are successful, deploy the container to your production environment and run acceptance tests.
D. 1. Alter the developers push the code to a central repository, trigger Cloud Build to build the application container. Then run unit tests
2. If unit tests are successful, deploy the container to a testing environment and run integration tests.

3. If the integration tests are successful, deploy the container to a production environment and run acceptance tests.

Answer: (SHOW ANSWER)

The most efficient pipeline follows the " Fail Fast " principle (Option A). You should run unit tests before the expensive and time-consuming process of building a container image. If the code logic is broken, the build should fail immediately. Once the image is built and stored, it is deployed to a Testing Environment for integration and acceptance tests to ensure it works within the system architecture. Finally, after deploying to Production , you run smoke tests (basic health checks) to ensure the live environment is stable. Options C and D are less ideal because they either suggest running acceptance tests after reaching production or skipping local validation steps, which increases the risk of deploying broken code to users.

NEW QUESTION: 38

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: (SHOW ANSWER)

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

NEW QUESTION: 39

You are designing the storage configuration for a critical Cloud Storage bucket. This bucket will store trade confirmation documents for a financial services company. These documents are immutable and serve as a legal record. Your design must meet two primary data protection requirements:

* The system must allow for the rapid recovery of any document that is accidentally deleted or maliciously overwritten.

* Once a document is written, it must be impossible for anyone, including administrators, to delete or modify it for seven years.

What should you do?

A. Enable Object Versioning on the bucket. Additionally, configure the bucket with a Bucket Lock and a retention policy of seven years.

B. Grant only the roles/storage.objectCreator IAM role to the application ' s service account. Remove all other admin permissions from the bucket.

C. Use Storage Transfer Service to create a daily replication job that synchronizes the bucket to a separate archive bucket in another region.

D. Configure an Object Lifecycle Management rule on the bucket to ensure all objects are retained for at least seven years before they can be deleted.

Answer: A (LEAVE A REPLY)

To satisfy strict regulatory compliance (such as SEC Rule 17a-4), you must implement WORM (Write Once, Read Many) storage. According to Cloud Storage documentation , this is achieved through a Retention Policy and Bucket Lock . Once a Bucket Lock is applied, the retention policy cannot be shortened or removed-even by a Project Owner or Administrator-ensuring the 7-year legal requirement is met.
To address the " rapid recovery " requirement for accidental deletions or overwrites, Object Versioning is essential. When versioning is enabled, deleting or overwriting an object creates a " noncurrent " version instead of permanently removing the data. This allows for immediate restoration. Option D (Lifecycle Management) is incorrect because it can only delete or change classes of objects after a time period; it does not prevent a user with sufficient permissions from deleting the object manually before that time. Option B (IAM) is insufficient because it does not protect against malicious administrators. Only the combination of Object Versioning and a locked Retention Policy provides the non-repudiable protection required for financial trade documents.

NEW QUESTION: 40

You are responsible for the Google Cloud environment in your company Multiple departments need access to their own projects and the members within each department will have the same project responsibilities You want to structure your Google Cloud environment for minimal maintenance and maximum overview of 1AM permissions as each department's projects start and end You want to follow Google-recommended practices What should you do?

Freecram.net

A. Create a Google Group per department and add all department members to their respective groups Create a folder per department and grant the respective group the required 1AM permissions at the folder level Add the projects under the respective folders

B. Grant all department members the required 1AM permissions for their respective projects

C. Create a Google Group per department and add all department members to their respective groups Grant each group the required I AM permissions for their respective projects

D. Create a folder per department and grant the respective members of the department the required 1AM permissions at the folder level. Structure all projects for each department under the respective folders

Answer: (SHOW ANSWER)

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

NEW QUESTION: 41
Your company has a Google Cloud project that uses BigQuery for data warehousing. The VPN tunnel between the on-premises environment and Google Cloud is configured with Cloud VPN. Your security team wants to avoid data exfiltration by malicious insiders, compromised code, and accidental oversharing. What should you do?
A. Configure VPC Service Controls and configure Private Google Access for on-premises hosts.
B. Create a service account, grant the BigQuery JobUser role and Storage Object Viewer role to the service account, and remove all other Identity and Access Management (1AM) access from the project.
C. Configure Private Google Access.
D. Configure Private Service Connect.
Answer: (SHOW ANSWER)
To secure data from exfiltration by malicious insiders, compromised code or accidental oversharing, we use VPC Service controls https://cloud.google.com/vpc-service-controls/docs/overview For private access options, connect to services in VPC networks we use private service endpoints or VPC network peering. https://cloud.google.com/vpc/docs/private-access-options#connect-services

NEW QUESTION: 42
You are the cloud architect for a team developing a new application that integrates with the Gemini API. You observe a sudden, massive spike in API costs and usage. An investigation reveals that a developer accidentally committed a service account key with Gemini API access to a public GitHub repository, and a malicious actor is using it. The on-call SRE immediately disabled the key, stopping the unauthorized usage. Your responsibility is to lead the post-incident review to prevent future credential leaks. What should you do?
A. Initiate a formal disciplinary review for the developer who committed the key. Mandate that all future pull requests must be manually approved by a security team member before being merged.
B. Immediately migrate all applications to use Workload Identity Federation to access Google Cloud APIs from GitHub Actions, eliminating all service account keys from your CI/CD process.
C. Write a report detailing the financial impact of the unauthorized use of the service account key. Work with the finance department to track the costs and create a policy to bill the responsible team for any future cost overruns caused by credential leaks.
D. Schedule a blameless postmortem with the engineering team to map the incident timeline.Collaboratively identify gaps in the development workflow, and assign action items to implement automated pre-commit hooks and a CI/CD pipeline stage for secret scanning.
Answer: (SHOW ANSWER)
Site Reliability Engineering (SRE) and DevOps best practices state that post-incident reviews must be blameless. Punitive actions create an environment of fear, which can lead to teams hiding mistakes and delaying incident resolution. A constructive postmortem focuses on uncovering systemic systemic failures within tools, workflows, and guardrails rather than assigning individual blame.
The primary goal of this review is to identity security gaps and implement automated preventative guardrails.

This involves adding secret detection tools (e.g., GitGuardian or TruffleHog) directly into development workflows using pre-commit hooks, as well as adding static analysis security testing (SAST) scanning stages into remote CI/CD deployment pipelines. This architecture catches exposed service keys or credentials before code changes are merged into public repositories, mitigating risk through automated technical controls.
References: Cloud Architect (GCP) Study Guide topics: " Google SRE Culture and Blameless Postmortems "
, " Secret Management Best Practices " , " Securing CI/CD Pipelines " , and " Credential Leak Mitigation " .

NEW QUESTION: 43
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: (SHOW ANSWER)

NEW QUESTION: 44
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()
```

freecram.net

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)
https://stackoverflow.c om/questi ons/3164280/google-app-engine-cache-list-in- session-variable?rq=1

NEW QUESTION: 45
The current Dress4win system architecture has high latency to some customers because it is located in one data center.
As of a future evaluation and optimizing for performance in the cloud, Dresss4win wants to distribute it's system architecture to multiple locations when Google cloud platform.
Which approach should they use?
A. Use regional managed instance groups and a global load balancer to increase reliability by providing automatic failover between zones in different regions.

B. Use a global load balancer with a set of virtual machines that forward the requests to a closer group of virtual machines managed by your operations team.
C. Use a global load balancer with a set of virtual machines that forward the requests to a closer group of virtual machines as part of a separate managed instance groups.
D. Use regional managed instance groups and a global load balancer to increase performance because the regional managed instance group can grow instances in each region separately based on traffic.
Answer: (SHOW ANSWER)

NEW QUESTION: 46
Your company is building containerized applications as pan of their CI ' CD pipeline. To improve the security and maintainability of the build process, you need to:
* Identify potential vulnerabilities within your container images.
* Generate verifiable metadata about the builds for auditing and compliance.
* Create a comprehensive inventory of your application ' s dependencies.
What should you do?
A. Use Cloud Build to build container images, and then trigger Artifact Analysis on images pushed to Artifact Registry.
B. Use Cloud Build to build container images, trigger Binary Authorization, and use Security Command Center for tracking and analysis
C. Use Cloud Build to build container images, push the images to Artifact Registry, and use Security Command Center for tracking and analysis
D. Use Cloud Build to build container images, trigger Binary Authorization, and use Cloud Asset Inventory for tracking and analysis.
Answer: (SHOW ANSWER)
The core of this question lies in three specific requirements: scanning for vulnerabilities, generating verifiable metadata, and creating a dependency inventory. Artifact Analysis (integrated with Artifact Registry) is the purpose-built Google Cloud service for these tasks. When an image is pushed, Artifact Analysis automatically performs a vulnerability scan against known CVE databases. It also generates build provenance (verifiable metadata) that proves the image was built by Cloud Build, which is essential for Binary Authorization to ensure only " trusted " images are deployed. Finally, it generates a Software Bill of Materials (SBOM) , which satisfies the requirement for a comprehensive inventory of dependencies. While Security Command Center (Option C) is a great dashboard for viewing risks, it is not the service that generates the metadata or performs the deep container analysis; it merely reports what Artifact Analysis finds.

Valid Professional-Cloud-Architect Dumps shared by ExamDiscuss.com for Helping Passing Professional-Cloud-Architect Exam! ExamDiscuss.com now offer the **newest Professional-Cloud-Architect exam dumps**, the ExamDiscuss.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Professional-

Cloud-Architect dumps with Test Engine here: https://www.examdiscuss.com/Google/exam/Professional-Cloud-Architect/premium/ (**353** Q&As Dumps, **35%OFF** **Special Discount Code:** **freecram**)

NEW QUESTION: 47
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 spot 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: 48
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)
https://cloud.google.com/functions/docs/securing/authenticating#authenticating_function_to_function_calls

NEW QUESTION: 49
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 add additional unit tests and production scale load tests on their cloud staging environment.
B. They should enable google stack driver debugger on the application code to show errors in the code
C. They should add canary tests so developers can measure how much of an impact the new release causes to latency

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

Answer: (SHOW ANSWER)

NEW QUESTION: 50

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

Freecram.net

A. Create a scalable environment in GCP for simulating production load.

B. Use the existing infrastructure to test the GCP-based backend at scale.

C. Build stress tests into each component of your application using resources internal to GCP to simulate load.

D. Create a set of static environments in GCP to test different levels of load - for example, high, medium, and low.

Answer: (SHOW ANSWER)

From scenario: Requirements for Game Backend Platform

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

Topic 3, JencoMart Case Study

Company Overview

JencoMart is a global retailer with over 10,000 stores in 16 countries. The stores carry a range of goods, such as groceries, tires, and jewelry. One of the company's core values is excellent customer service. In addition, they recently introduced an environmental policy to reduce their carbon output by 50% over the next 5 years.

Company Background

JencoMart started as a general store in 1931, and has grown into one of the world's leading brands known for great value and customer service. Over time, the company transitioned from only physical stores to a stores and online hybrid model, with 25% of sales online. Currently, JencoMart has little presence in Asia, but considers that market key for future growth.

Solution Concept

JencoMart wants to migrate several critical applications to the cloud but has not completed a technical review to determine their suitability for the cloud and the engineering required for migration. They currently host all of these applications on infrastructure that is at its end of life and is no longer supported.

Existing Technical Environment

JencoMart hosts all of its applications in 4 data centers: 3 in North American and 1 in Europe, most applications are dual-homed.

JencoMart understands the dependencies and resource usage metrics of their on-premises architecture.

Application Customer loyalty portal

LAMP (Linux, Apache, MySQL and PHP) application served from the two JencoMart-owned U.S. data centers.

Database

* Oracle Database stores user profiles
* 20 TB
* Complex table structure
* Well maintained, clean data
* Strong backup strategy
* PostgreSQL database stores user credentials
* Single-homed in US West

No redundancy

Backed up every 12 hours

* 100% uptime service level agreement (SLA)
* Authenticates all users

Compute

* 30 machines in US West Coast, each machine has:

Twin, dual core CPUs

32GB of RAM

* Twin 250 GB HDD (RAID 1)
* 20 machines in US East Coast, each machine has:

Single dual-core CPU

24 GB of RAM

* Twin 250 GB HDD (RAID 1)

Storage

* Access to shared 100 TB SAN in each location
* Tape backup every week

Business Requirements

* Optimize for capacity during peak periods and value during off-peak periods
* Guarantee service availably and support
* Reduce on-premises footprint and associated financial and environmental impact.
* Move to outsourcing model to avoid large upfront costs associated with infrastructure purchase
* Expand services into Asia.

Technical Requirements

* Assess key application for cloud suitability.
* Modify application for the cloud.
* Move applications to a new infrastructure.
* Leverage managed services wherever feasible
* Sunset 20% of capacity in existing data centers
* Decrease latency in Asia

CEO Statement

JencoMart will continue to develop personal relationships with our customers as more people access the web.
The future of our retail business is in the global market and the connection between online and in-store experiences. As a large global company, we also have a responsibility to the environment through 'green' initiatives and polices.
CTO Statement
The challenges of operating data centers prevents focus on key technologies critical to our long-term success.
Migrating our data services to a public cloud infrastructure will allow us to focus on big data and machine learning to improve our service customers.
CFO Statement
Since its founding JencoMart has invested heavily in our data services infrastructure. However, because of changing market trends, we need to outsource our infrastructure to ensure our long-term success. This model will allow us to respond to increasing customer demand during peak and reduce costs.

NEW QUESTION: 51
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. Configure the deployment job to notify a Pub/Sub queue that triggers a Cloud Function.
B. Set up a Cloud Logging sink and a Cloud Storage bucket that triggers a Cloud Function.
C. Set up Cloud Tasks and a Cloud Storage bucket that triggers a Cloud Function.
D. Set up Identity and Access Management (IAM) and Confidential Computing to trigger a Cloud Function.
Answer: (SHOW ANSWER)

NEW QUESTION: 52
You have an application that runs in Google Kubernetes Engine (GKE). Over the last 2 weeks, customers have reported that a specific part of the application returns errors very frequently. You currently have no logging or monitoring solution enabled on your GKE cluster. You want to diagnose the problem, but you have not been able to replicate the issue. You want to cause minimal disruption to the application. What should you do?
A. 1. Update your GKE cluster to use Cloud Operations for GKE.
2. Use the GKE Monitoring dashboard to investigate logs from affected Pods.
B. 1. Create a new GKE cluster with Cloud Operations for GKE enabled.
2.Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluster.
3.Use the GKE Monitoring dashboard to investigate logs from affected Pods.
C. 1. Update your GKE cluster to use Cloud Operations for GKE, and deploy Prometheus.
2. Set an alert to trigger whenever the application returns an error.

D. 1. Create a new GKE cluster with Cloud Operations for GKE enabled, and deploy Prometheus.
2.Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluster.
3.Set an alert to trigger whenever the application returns an error.
Answer: (SHOW ANSWER)
Reference: https://cloud.google.com/blog/products/management-tools/using-logging-your-apps-running-kubernetes-engine

NEW QUESTION: 53
To reduce costs, the Director of Engineering has required all developers to move their development infrastructure resources from on-premises virtual machines (VMs) to Google Cloud Platform. These resources go through multiple start/stop events during the day and require state to persist. You have been asked to design the process of running a development environment in Google Cloud while providing cost visibility to the finance department. Which two steps should you take? Choose 2 answers
A. Use the --no-auto-delete flag on all persistent disks and stop the VM.
B. Use the -auto-delete flag on all persistent disks and terminate the VM.
C. Apply VM CPU utilization label and include it in the BigQuery billing export.
D. Use Google BigQuery billing export and labels to associate cost to groups.
E. Store all state into local SSD, snapshot the persistent disks, and terminate the VM.
F. Store all state in Google Cloud Storage, snapshot the persistent disks, and terminate the VM.
Answer: (SHOW ANSWER)
https://cloud.google.com/billing/docs/how-to/export-data-bigquery

NEW QUESTION: 54
You are developing a retail application on Vertex AI that uses a large language model to give customers real- time product recommendations. The app uses text, code, images, audio, and video prompts to communicate with customers. You need to minimize the application ' s latency to ensure a responsive and delightful user experience. What should you do?
A. Enable streaming responses from the model to process the output in real time.
B. Increase the temperature parameter.
C. Avoid using the max_output_tokens parameter to prevent overly long responses.
D. Avoid using system instructions to prevent overly long responses.
Answer: (SHOW ANSWER)
According to Google Cloud Vertex AI documentation on Generative AI , Streaming is the recommended technique to reduce perceived latency in user-facing applications. When using large language models (LLMs), generating the full response can take several seconds depending on the complexity and length. By enabling streaming, the model sends the output in chunks (tokens) as they are generated, rather than waiting for the entire sequence to complete. This allows the application UI to display text to the customer progressively, creating a " real-time " feel and a more responsive experience.
Options B, C, and D do not directly address latency in a reliable or architectural way. Increasing temperature (Option B) affects the randomness and creativity of the output but not the speed of generation. While setting a lower max_output_tokens (Option C) might technically result in shorter (and thus faster)

responses, " avoiding " the parameter altogether gives the model more freedom to produce long, slow responses, which is counterproductive. Similarly, avoiding system instructions (Option D) sacrifices the quality and safety of the model ' s behavior without providing a predictable latency benefit. For a high-quality " conversational commerce " experience as defined in the Cymbal Retail case study, streaming is the standard architectural pattern for performance.

NEW QUESTION: 55
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: (SHOW ANSWER)
Reference https://cloud.google.com/load-balancing/docs/https/url-map

NEW QUESTION: 56
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. Use regular Compute Engine instances for the production environment, and use Spot VMs for the acceptance and development environments.
B. Deploy the development and acceptance applications on a managed instance group and enable autoscaling.
C. 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.
D. Use Instance Schedules to stop the development and acceptance environments after office hours, and start them just before office hours.
Answer: D (LEAVE A REPLY)
Comprehensive and Detailed Explanation From Exact Extract: The most effective way to achieve significant cost savings for non-production environments that are only needed during business hours is to stop the instances when they are idle. When an instance is stopped, you are only charged for the persistent disks and not for the CPU/RAM.
Instance Schedules are the native, managed Google Cloud feature designed specifically for this purpose: automating the starting and stopping of Compute Engine instances on a recurring schedule.

* Option D (Use Instance Schedules to stop the development and acceptance environments after office hours, and start them just before office hours) directly addresses the need for cost optimization during idle times with the recommended native solution.
* Option A (Spot VMs) are cheap but can be preempted at any time, which isn ' t suitable for critical during office hours environments.
* Option B (Autoscaling) scales down when demand is low, but typically won ' t scale completely to zero instances, and will still incur minimum running costs.
* Option C (Custom Shell Script) is a manual, non-native solution that requires maintenance and is more complex than the built-in Instance Schedules.
Exact Extract Principle:
From the Google Cloud documentation on Compute Engine Cost Optimization :
" Instance schedules allow you to schedule when Compute Engine virtual machine instances start and stop. Use instance schedules to automatically start and stop your instances at predefined times. This is a highly recommended practice for cost optimization of non-production environments (development, testing, or staging) that are not needed 24/7, as you are not billed for the instance ' s machine type when it is stopped.
" Reference: Google Cloud Compute Engine Instance Schedules Documentation and Compute Engine Pricing
/Cost Optimization Guide.

NEW QUESTION: 57
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: (SHOW ANSWER)
https://cloud.google.com/transfer-appliance/docs/2.0/faq

NEW QUESTION: 58
Your company provides a recommendation engine for retail customers. You are providing retail customers with an API where they can submit a user ID and the API returns a list of recommendations for that user. You are responsible for the API lifecycle and want to ensure stability for your customers in case the API makes backward-incompatible changes. You want to follow Google-recommended practices. What should you do?
A. Create a distribution list of all customers to inform them of an upcoming backward-incompatible change at least one month before replacing the old API with the new API.

B. Create an automated process to generate API documentation, and update the public API documentation as part of the CI/CD process when deploying an update to the API.
C. Use a versioning strategy for the APIs that increases the version number on every backward-incompatible change.
D. Use a versioning strategy for the APIs that adds the suffix "DEPRECATED" to the current API version number on every backward-incompatible change. Use the current version number for the new API.
Answer: (SHOW ANSWER)
https://cloud.google.com/apis/design/versioning
All Google API interfaces must provide a major version number, which is encoded at the end of the protobuf package, and included as the first part of the URI path for REST APIs. If an API introduces a breaking change, such as removing or renaming a field, it must increment its API version number to ensure that existing user code does not suddenly break.

NEW QUESTION: 59
For this question, refer to the Mountkirk Games case study. Mountkirk Games wants you to design a way to test the analytics platform's resilience to changes in mobile network latency. What should you do?
A. Create an opt-in beta of the game that runs on players' mobile devices and collects response times from analytics endpoints running in Google Cloud Platform regions all over the world.
B. Add the ability to introduce a random amount of delay before beginning to process analytics files uploaded from mobile devices.
C. Deploy failure injection software to the game analytics platform that can inject additional latency to mobile client analytics traffic.
D. Build a test client that can be run from a mobile phone emulator on a Compute Engine virtual machine, and run multiple copies in Google Cloud Platform regions all over the world to generate realistic traffic.
Answer: (SHOW ANSWER)

NEW QUESTION: 60
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?
Freecram.net
A. Enable autoscaling on the GKE cluster and enable horizontal pod autoscaling on your application deployments. Send curl requests to your application, and validate if the auto scaling works.
B. Use a load testing tool to simulate the expected number of concurrent users and total requests to your application, and inspect the results.
C. Replicate the application over multiple GKE clusters in every Google Cloud region. Configure a global HTTP (S) load balancer to expose the different clusters over a single global IP address.
D. Use Cloud Debugger in the development environment to understand the latency between the different microservices.
Answer: (SHOW ANSWER)

NEW QUESTION: 61
You are migrating third-party applications from optimized on-premises virtual machines to Google Cloud. You are unsure about the optimal CPU and memory requirements. The applications have a con sistent usage pattern across multiple weeks. You want to optimize resource usage for the lowest cost . What should you do?
A. 1. Create multiple Compute Engine instances with varying CPU and memory options.
2. Install the Ops Agent, and deploy the third-party application on each of them.
3. Run a load test with high traffic levels on the application, and use the results to determine the optimal settings.
B. 1. Create an instance template with the smallest available machine type, and use an image of the third-party application taken from a current on-premises virtual machine.
2. Create a managed instance group that uses average CPU utilization to autoscale the number of instances in the group.
3. Modify the average CPU utilization threshold to optimize the number of instances running.
C. 1. Create an App Engine flexible environment, and deploy the third-party application using a Dockerfile and a custom runtime.
2. Set CPU and memory options similar to your application ' s current on-premises virtual machine in the app.yaml file.
D. 1. Create a Compute Engine instance with CPU and memory options similar to your application's current on-premises virtual machine.
2. Install the Ops agent, and deploy the third-party application.
3. Run a load test with normal traffic levels on the application, and follow the Rightsizing Recommendations in the Google Cloud console.
Answer: (SHOW ANSWER)
Comprehensive and Detailed Explanation From Exact Extract: The goal is to optimize resource usage for the lowest cost when the optimal CPU and memory are unknown, and the application has a consistent usage pattern .
The best Google Cloud-native approach for this is to leverage Rightsizing Recommendations . This feature analyzes the historical performance data of your VMs to suggest a machine type that is more closely aligned with actual usage.
* Step 1 provides a reasonable starting point.
* Step 2 (Install the Ops agent) is crucial as it collects detailed system-level metrics (CPU, memory, disk, network) required for generating accurate rightsizing recommendations.
* Step 3 (Run a load test with normal traffic levels and follow the Rightsizing Recommendations) ensures that the observed usage is based on real-world (or simulated normal-world) load and uses the automated Google Cloud tool to find the most cost-effective machine type. Running a load test is necessary to generate sufficient, representative telemetry data.
* Option A ' s step 3 uses high traffic, which would lead to an oversized instance recommendation.
* Option B and C don ' t leverage the Rightsizing Recommendations feature, which is the most efficient and automated way to find optimal resource usage based on actual data.
Exact Extract Principle:

From the Google Cloud documentation on Rightsizing Recommendations :
" Google Cloud provides Rightsizing Recommendations for Compute Engine instances based on resource usage metrics collected by the Ops Agent or legacy monitoring agent. These recommendations help you optimize your resources for cost and performance. To generate accurate and effective rightsizing recommendations , the instance must be running a representative workload and have the Ops Agent installed to collect CPU and memory utilization data. "
" Rightsizing recommendations are a key feature for cost optimization in Compute Engine, suggesting smaller or larger machine types that match your actual utilization. " Reference: Google Cloud Compute Engine Rightsizing Recommendations Documentation.

Valid Professional-Cloud-Architect Dumps shared by ExamDiscuss.com for Helping Passing Professional-Cloud-Architect Exam! ExamDiscuss.com now offer the **newest Professional-Cloud-Architect exam dumps**, the ExamDiscuss.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Professional-Cloud-Architect dumps with Test Engine here: https://www.examdiscuss.com/Google/exam/Professional-Cloud-Architect/premium/ (**353** Q&As Dumps, **35%OFF** **Special Discount Code: freecram**)

NEW QUESTION: 62
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: (SHOW ANSWER)
https://cloud.google.com/storage/docs/gsutil/commands/lifecycle

NEW QUESTION: 63
You are helping the QA team to roll out a new load-testing tool to test the scalability of your primary cloud services that run on Google Compute Engine with Cloud Bigtable. Which three requirements should they include? Choose 3 answers
A. Instrument the production services to record every transaction for replay by the load-testing tool.
B. Schedule the load-testing tool to regularly run against the production environment.
C. Ensure all third-party systems your services use are capable of handling high load.
D. Instrument the load-testing tool and the target services with detailed logging and metrics collection.
E. Create a separate Google Cloud project to use for the load-testing environment.
F. Ensure that the load tests validate the performance of Cloud Bigtable.
Answer: D,E,F (LEAVE A REPLY)

NEW QUESTION: 64
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: (SHOW ANSWER)
https://stackoverflow.com/questions/36441423/migrate-google-compute-engine-instance-to-a-different-region

NEW QUESTION: 65
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: (SHOW ANSWER)
https://cloud.google.com/solutions/federating-gcp-with-active-directory-introduction#implementing_federation

NEW QUESTION: 66
Your company just finished a rapid lift and shift to Google Compute Engine for your compute needs. You have another 9 months to design and deploy a more cloud-native solution. Specifically, you want a system that is no-ops and auto-scaling. Which two compute products should you choose? Choose 2 answers
A. Compute Engine with containers
B. Google Kubernetes Engine with containers
C. Google App Engine Standard Environment
D. Compute Engine with custom instance types
E. Compute Engine with managed instance groups
Answer: B,C (LEAVE A REPLY)

B: With Container Engine, Google will automatically deploy your cluster for you, update, patch, secure the nodes.
Kubernetes Engine's cluster autoscaler automatically resizes clusters based on the demands of the workloads you want to run.
C: Solutions like Datastore, BigQuery, AppEngine, etc are truly NoOps.
App Engine by default scales the number of instances running up and down to match the load, thus providing consistent performance for your app at all times while minimizing idle instances and thus reducing cost.
Note: At a high level, NoOps means that there is no infrastructure to build out and manage during usage of the platform. Typically, the compromise you make with NoOps is that you lose control of the underlying infrastructure.
References: https://www.quora.com/How-well-does-Google-Container-Engine-support-Google-Cloud-Platform%E2%80%99s-NoOps-claim

NEW QUESTION: 67
You are migrating a critical, on-premises, three-tier ecommerce application to Google Cloud. The application has tightly coupled dependencies between its web, application, and database tiers. The business has strict performance and availability requirements and is concerned about the risks associated with an " all-at-once " migration. You want to follow Google-recommended practices to create a migration plan that systematically addresses application dependencies and validates post-migration performance before the final production cutover What should you do?
A. Redesign the application to be cloud-native using microservices before migration to eliminate dependencies and simplify workload testing.
B. Plan a " lift and shift " migration for all application tiers to be executed simultaneously over a single weekend to minimize downtime.
C. Use dependency mapping to group application components into waves, create a phased migration plan, and incorporate performance testing for each wave.
D. Prioritize the establishment of a high-bandwidth, low-latency Cloud Interconnect connection, and then plan to migrate the database servers first.
Answer: (SHOW ANSWER)
According to the Google Cloud Migration Framework , the recommended approach for migrating complex, tightly coupled applications is to follow a structured process: Discover, Assess, Plan, and Migrate . For high- risk environments, a " Big Bang " migration (Option B) is discouraged because it provides no room for validation and creates a high probability of failure.
By using dependency mapping (often via tools like Migration Center), Altostrat can identify which components must move together. Grouping these into " waves " allows for a phased migration . This strategy mitigates risk by moving non-critical or less complex parts first to build confidence. Most importantly, it allows for rigorous performance testing after each wave to ensure that the hybrid connectivity and cloud infrastructure meet the strict ecommerce performance requirements before the final cutover.
Option A (refactoring) adds too much time and risk before the migration even begins, and Option D ignores the application-layer dependencies that could break if the database is moved in isolation.

NEW QUESTION: 68
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: (SHOW ANSWER)
https://cloud.google.com/endpoints/docs/openapi/about-cloud-endpoints?hl=en_US & _ga=2.21 787131. -1712523161.1522785064
https:/ /cloud.google.com/endpoints/docs/openapi/architecture-overview
https://cloud.google.com/storage/docs/gsutil/commands/test
Develop, deploy, protect and monitor your APIs with Google Cloud Endpoints. Using an Open API Specification or one of our API frameworks, Cloud Endpoints gives you the tools you need for every phase of API development.
From scenario:
Business Requirements
Decrease unplanned vehicle downtime to less than 1 week, without increasing the cost of carrying surplus inventory Support the dealer network with more data on how their customers use their equipment to better position new products and services Have the ability to partner with different companies - especially with seed and fertilizer suppliers in the fast- growing agricultural business - to create compelling joint offerings for their customers.
Reference: https://cloud.google.com/certification/guides/cloud-architect/casestudy-terramearth

NEW QUESTION: 69
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: (SHOW ANSWER)
https://cloud.google.com/interconnect/docs/how-to/direct-peering

NEW QUESTION: 70
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?
Freecram.net
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: 71
You are a cloud architect for a financial services organization that is migrating its environment to Google Cloud. A new developer on your team needs to quickly inspect the metadata of a Compute Engine instance to troubleshoot a configuration issue, but they do not have the Google Cloud SDK installed on their local workstation. You need to provide a quick and effective solution that allows them to run gcloud commands directly from their browser. What should you do?
A. Instruct the developer to install Cloud Code in their local integrated development environment (IDE).
B. Tell the developer to provision a Cloud Workstation and use its terminal to run the commands.
C. Guide the developer to download and install the Google Cloud SDK on their workstation.
D. Advise the developer to use the Cloud Shell Terminal.
Answer: (SHOW ANSWER)
The Google Cloud Console provides Cloud Shell as an embedded, browser-accessible terminal environment that removes local machine installation overhead. Cloud Shell comes pre-provisioned with the Google Cloud SDK (gcloud CLI), command-line tools, and pre-authenticated context based on the user ' s

active console session. This fulfills the strict requirement to quickly run management operations directly from a web browser without executing local software installers or modifying the workstation.
While Cloud Workstations offers a fully managed development environment in the cloud, provisioning an entirely new instance introduces unnecessary infrastructure provisioning delays and operational complexity for a simple metadata inspection task. Installing Cloud Code or downloading the SDK locally directly violates the constraints of avoiding local workstation modifications. Cloud Shell provides a temporary compute instance backed by an ephemeral container, making it the most immediate, frictionless utility for running ad- hoc operations.
References: Cloud Architect (GCP) Study Guide topics: " Google Cloud Management Tools " , " Cloud Shell Utilities and Ephemeral Storage " , and " Interacting with Compute Engine Metadata via gcloud " .

NEW QUESTION: 72
Your company runs a critical, revenue-generating ecommerce application that is served by a regional managed instance group (MIG) behind an external HTTP(S) Load Balancer. The operations team is currently overwhelmed with low-priority notifications. Your team ' s service level objective (SLO) is to maintain 99.9% availability. You want to minimize noise from non-critical events and ensure that the team is only notified of issues that are actionable and threaten the SLO. What should you do?
A. Focus on cause-based alerts, creating alerting policies with thresholds for the Compute Engine instances, including CPU utilization, memory usage, disk I/O, and network traffic.
B. Configure alerts based on predictive metrics. Use the instance count of the MIG as the primary metric to trigger an alert.
C. Implement an error budget policy based on the availability of the SLO. Create a " page " alert that triggers only when the rate of burn of the error budget predicts a full exhaustion within the next 24 hours.
D. Create log-based alerts for only the WARN and ERROR log entries generated by the application to ensure that no potential issue is missed.
Answer: (SHOW ANSWER)
According to the Google SRE Handbook , modern monitoring should focus on Symptom-based alerting rather than Cause-based alerting. Option A (Cause-based) leads to " alert fatigue " because high CPU or memory usage may not actually be impacting the user experience if the application is still responding successfully.
The most effective way to protect a 99.9% availability SLO is to use Error Budgets and Burn Rate Alerting . An error budget is the amount of unreliability your service can tolerate (0.1% for a 99.9% SLO).
By implementing an alert that triggers when the burn rate is high enough to exhaust the budget within 24 hours (Option C), the team is only " paged " for significant, persistent issues that actually threaten the business goal. Smaller, momentary spikes that don ' t meaningfully impact the long-term SLO can be addressed during normal business hours as " low-priority tickets. " This approach significantly reduces noise, improves the signal- to-noise ratio, and ensures that when the team is notified, the issue is critical and requires immediate action to preserve revenue and customer trust.

NEW QUESTION: 73

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

A. Apply a Cloud Armor security policy to external load balancers using a named IP list for Fastly.

B. Apply a Cloud Armor security policy to external load balancers using the IP addresses that Fastly has published. C. Apply a VPC firewall rule on port 443 for Fastly IP address ranges.

C. Apply a VPC firewall rule on port 443 for network resources tagged with scurceiplisr-fasrly.

Answer: (SHOW ANSWER)

Reference: https://cloud.google.com/load-balancing/docs/https

D18912E1457D5D1DDCBD40AB3BF70D5D

NEW QUESTION: 74

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

What should you do?

A. Use Stackdriver Trace to create a trace list analysis.

B. Use Stackdriver Monitoring to create a dashboard on the project's activity.

C. Enable Cloud Identity-Aware Proxy in all projects, and add the group of Administrators as a member.

D. Use the Activity page in the GCP Console and Stackdriver Logging to provide the required insight.

Answer: (SHOW ANSWER)

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

NEW QUESTION: 75

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

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

A. Tests should scale well beyond the prior approaches.

B. Unit tests are no longer required, only end-to-end tests.

C. Tests should be applied after the release is in the production environment.

D. Tests should include directly testing the Google Cloud Platform (GCP) infrastructure.

Answer: (SHOW ANSWER)

From Scenario:

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

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

NEW QUESTION: 76
You are setting up a new regional environment for your business needs. You plan to migrate more than 2 TB of data from the us-central1 region to the asia-southeast1 region. The transfer should occur as quickly as possible. If the transfer is interrupted, the operation must resume where it left off and copy only new or modified objects. What should you do?
A. Utilize Storage Transfer Service (STS) to migrate objects to the destination bucket in the target region. Once the transfer is complete, delete the original source bucket objects.
B. Use the gcloud storage rsync --recursive command to replicate the source bucket to a new target bucket in the us-west1 region.
C. Use the gcloud storage cp command to copy the objects. Be careful to manage potential naming conflicts to ensure data integrity.
D. Leverage Cloud Composer to orchestrate a data pipeline that copies objects from the source bucket to the destination bucket.
Answer: (SHOW ANSWER)
According to Google Cloud Storage documentation , the Storage Transfer Service (STS) is the purpose-built tool for large-scale data movement between Cloud Storage buckets, especially across regions.
STS is superior for this scenario for several reasons. First, it is a fully managed service that operates on Google's backbone, providing much higher throughput and reliability than a client-side command like gcloud storage cp (Option C). Second, STS inherently supports incremental transfers and checkpointing . If a transfer job is interrupted, it automatically resumes and only transfers objects that haven ' t been successfully copied or have been modified since the last run, satisfying the specific requirement of the prompt.
While gcloud storage rsync (Option B) provides similar synchronization logic, it runs on the user ' s local machine or a specific VM, making it subject to the performance and stability of that single instance. STS, being a managed service, scales out resources to handle 2 TB+ transfers " as quickly as possible. " Option D (Cloud Composer) adds unnecessary complexity and cost for a task that is natively handled by STS. Option A provides the most robust, performant, and automated path for regional data migration.

Valid Professional-Cloud-Architect Dumps shared by ExamDiscuss.com for Helping Passing Professional-Cloud-Architect Exam! ExamDiscuss.com now offer the **newest Professional-Cloud-Architect exam dumps**, the ExamDiscuss.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Professional-Cloud-Architect dumps with Test Engine here: https://www.examdiscuss.com/Google/exam/Professional-Cloud-Architect/premium/ (**353** Q&As Dumps, **35%OFF** **Special Discount Code:** **freecram**)

NEW QUESTION: 77
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: (SHOW ANSWER)
https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer
https://cloud.google.com/kubernetes-engine/docs/concepts/network-overview#ext-lb

NEW QUESTION: 78
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 on-premises database.
Answer: (SHOW ANSWER)
https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databases

NEW QUESTION: 79
You have a Compute Engine application that you want to autoscale when total memory usage exceeds 80%.
You have Installed the Cloud Ops agent and configured the autoscaling policy as follows:
You observe that the application does not scale under high load. You want to resolve this. What should you do?
A. Change the Target type to DELTA_PER_MINUTE
B. Change the metric identifier to agent.googleapis.com/memory/bytes_used
C. Change the filer to metric.label.state = 'free' and the Target utilization level to 20
D. Change the Filter to metric.label.state = 'used'
Answer: (SHOW ANSWER)

NEW QUESTION: 80

A recent audit that a new network was created in Your GCP project. In this network, a GCE instance has an SSH port open the world. You want to discover this network's origin. What should you do?

Freecram.net

A. Navigate to the Activity page in the Home section. Set category to Data Access and search for Create VM entry.

B. Connect to the GCE instance using project SSH Keys. Identify previous logins in system logs, and match these with the project owners list.

C. Search for Create VM entry in the Stackdriver alerting console.

D. In the logging section of the console, specify GCE Network as the logging section. Search for the Create Insert entry.

Answer: (SHOW ANSWER)

NEW QUESTION: 81

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

A. Help the engineer to convert his websocket code to use HTTP streaming.

B. Review the encryption requirements for websocket connections with the security team.

C. Meet with the cloud operations team and the engineer to discuss load balancer options.

D. Help the engineer redesign the application to use a distributed user session service that does not rely on websockets and HTTP sessions.

Answer: C (LEAVE A REPLY)

Google Cloud Platform (GCP) HTTP(S) load balancing provides global load balancing for HTTP(S) requests destined for your instances.

The HTTP(S) load balancer has native support for the WebSocket protocol.

References: https://cloud.google.com/compute/docs/load-balancing/http/

NEW QUESTION: 82

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)

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_migration

NEW QUESTION: 83

Your company is running AI/ML workloads on graphics processing units (GPUs) within Google Kubernetes Engine (GKE). These workloads are deployed across production and development projects, with the production environment being business critical and operating 24/7. You need to ensure that the GPU capacity is always available during scaling events, without being limited by overall GPU availability. You also need the ability to shift unused GPU capacity from production to development projects as needed. You want to follow Google-recommended best practices. What should you do?

A. Create a reservation in the production project, and then create a separate reservation in the development project to secure capacity.

B. Create a reservation in the production project, and then transfer the reservation to the development project when needed.

C. Create an owner project with shared reservations, and then configure the production project and development project as consumer projects to consume capacity.

D. Create a shared reservation in the production project, and then add the development project as a consumer project to move capacity when needed.

Answer: (SHOW ANSWER)

According to Compute Engine documentation on Shared Reservations , the most flexible and manageable way to handle capacity across multiple projects is to use a centralized Owner Project . This is a dedicated project used to create and manage the reservation, which is then shared with " Consumer Projects " (in this case, your GKE production and development projects).

This architecture specifically addresses the requirement to " ensure GPU capacity is always available " by securing the resources at the organizational level. Because the reservation is shared , the GPU capacity is not

" locked " into a single project. If the production environment scales down and leaves unused capacity in the shared pool, the development project can immediately consume those reserved GPUs for its own scaling needs without requiring manual transfers or new configurations.

Google ' s Best Practices for Shared Reservations recommend that the owner project should only own the reservation and not necessarily consume it, which simplifies billing and quota management. Option D is technically possible but less " manageable " at scale than a centralized owner project. Options A and B lack the automated, fluid capacity shifting required by the AI/ML workload ' s dynamic scaling needs.

The following solutions are based on Google Cloud Architect documents and industry-standard best practices for global content delivery, data migration, and high-performance file access.

NEW QUESTION: 84

You are designing me observability strategy for a new microservices application running on Google Kubernetes Engine (GKE) The application consists of multiple services (e.g.. frontend, orders, payments). During load testing, you observe an error in the frontend service ' s logs, but you cannot find the corresponding logs in the downstream services to investigate the root cause because the logs are not correlated. You need to implement a solution that allows you to follow a single user request across all microservices involved in the transaction. The solution must not require developers to manually add correlation logic to their application code What should you do?

A. Require developers to generate a unique correlation-id at the frontend. and manually add it as a field to every log message in all services.
B. Configure all containers to write logs to STDOUTSTDERR. and then filter logs by pod name and timestamp m the Logs Explorer
C. Implement Cloud Trace by ensuring the traceparent header is propagated between microservice calls to link logs to a single trace
D. Create custom metrics in Cloud Monitoring for error counts in each service and correlate incident spikes using a shared dashboard
Answer: (SHOW ANSWER)
To follow a user request through a maze of microservices, you need Distributed Tracing .
* Cloud Trace: By ensuring the traceparent (W3C standard) header is propagated (Option C), Google Cloud can automatically link logs from the Frontend, Orders, and Payments services.
* Zero Manual Effort: Modern libraries (like OpenTelemetry) handle the " heavy lifting " of extracting and injecting these headers into outgoing requests, satisfying the requirement to avoid manual correlation logic.

NEW QUESTION: 85
You need to upgrade the EHR connection to comply with their requirements. The new connection design must support business-critical needs and meet the same network and security policy requirements. What should you do?
A. Add a new Carrier Peering connection.
B. Upgrade the bandwidth on the Dedicated Interconnect connection to 100 G.
C. Add three new Cloud VPN connections.
D. Add a new Dedicated Interconnect connection.
Answer: (SHOW ANSWER)

NEW QUESTION: 86
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)
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: 87
You are implementing Firestore for Mountkirk Games. Mountkirk Games wants to give a new game programmatic access to a legacy game ' s Firestore database. Access should be as restricted as possible. What should you do?
A. Create a service account (SA) in the legacy game ' s Google Cloud project, give it the Firebase Admin role, and then migrate the new game to the legacy game ' s project.
B. Create a service account (SA) in the legacy game ' s Google Cloud project, add this SA in the new game ' s IAM page, and then give it the Firebase Admin role in both projects
C. Create a service account (SA) in the legacy game ' s Google Cloud project, add a second SA in the new game ' s IAM page, and then give the Organization Admin role to both SAs

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

NEW QUESTION: 88
You are developing a deep learning model that requires high-performance access to large volumes of media data currently stored in Cloud Storage. Model training will be executed on multiple VM instances with GPUs attached, but your application must interact with the data as if it were on the local file system. You need to minimize complexity and cost. What should you do?
A. Use the gsutil command line tool to download the data to your VM instances.
B. Copy the data from Cloud Storage to Filestore, and then mount the Filestore volume as a local file system on your VM instances.
C. Create a shared persistent disk, attach the disk to your VM instances, and load data from the Cloud Storage bucket.
D. Use the gcsfuse command line tool to mount the Cloud Storage bucket as a local file system, and perform read/write operations in your bucket using standard file system semantics.
Answer: (SHOW ANSWER)
According to Google Cloud documentation on Cloud Storage FUSE , gcsfuse is the recommended tool for applications that require a POSIX-compliant file system interface to Cloud Storage buckets. This allows deep learning frameworks (like TensorFlow or PyTorch) to read media data using standard file I/O commands (open, read) as if the data were on a local disk, without requiring the application code to be rewritten to use the Cloud Storage API.
Cloud Storage FUSE is the most cost-effective and least complex solution because it avoids the need to duplicate data. Option B (Filestore) and Option C (Persistent Disk) would require you to pay for expensive block/file storage in addition to your Cloud Storage costs, and would require time-consuming data copying steps. Option A (download ing) is impractical for " large volumes of media data " as it would consume significant time and local disk space on every GPU instance.
While gcsfuse does not provide the same performance as a local SSD for random small-file I/O, it is highly optimized for the sequential read patterns typical of media processing and deep learning training. By mounting the bucket directly, you minimize operational overhead and only pay for the data you store in Cloud Storage and the egress associated with the reads, aligning perfectly with the goal of " minimizing complexity and cost. "

NEW QUESTION: 89
Altostrat ' s development team is using a microservices architecture for their application. You need to select the most suitable testing approach to ensure that individual microservices function correctly in isolation. What should you do?
A. Run unit testing.
B. Use load testing.
C. Perform end-to-end testing.
D. Execute integration testing.

Answer: (SHOW ANSWER)

According to the Google Cloud Architecture Framework: Software Development and Delivery section, unit testing is the foundational practice for verifying the smallest parts of an application-such as functions, classes, or individual microservices-in complete isolation . The requirement specifically asks for a method to ensure services function correctly " in isolation, " which is the primary definition of a unit test.

In Altostrat's modernized CI/CD environment, unit tests are executed early in the pipeline (the " left " side of DevSecOps). They typically use " mocks " or " stubs " to simulate external dependencies like BigQuery, Cloud Storage, or other microservices. This ensures that the internal logic of a specific service-for instance, the metadata extraction logic or the pricing algorithm-is mathematically and logically sound before it interacts with the rest of the system.

While Integration testing (Option D) is vital for ensuring services talk to each other correctly, it does not test them " in isolation. " End-to-end testing (Option C) validates the entire user journey but is slow and complex to debug. Load testing (Option B) focuses on performance under stress rather than functional correctness.

To meet the technical requirement of " Modernize CI/CD for containerized deployments, " Altostrat must prioritize a robust suite of unit tests to provide developers with rapid feedback and ensure the high reliability of their operational workflows.

NEW QUESTION: 90

You have an application that uses Vertex AI Feature Store to manage and serve product features for real-time recommendations. You want to monitor the performance and health of the application. You need to understand the overall duration of a request. What should you do?

Freecram.net

A. Measure the Latency of your requests.

B. Track the Online serving throughput of your requests.

C. Observe the Request size in your featurestore.

D. Monitor the Queries per second for your featurestore.

Answer: (SHOW ANSWER)

According to the Google Cloud SRE (Site Reliability Engineering) handbook and Vertex AI Monitoring documentation , Latency is the fundamental metric used to measure the time it takes to service a request. In the context of " real-time recommendations " for Cymbal Retail, understanding the overall duration -from the moment the app requests a feature to the moment it receives the response-is critical for maintaining the user experience.

In Vertex AI Feature Store, latency is typically measured at various percentiles (e.g., p50, p95, and p99) to account for outliers. High latency in the online serving path directly impacts the speed of the recommendation engine.

While Online serving throughput (Option B) and Queries per second (QPS) (Option D) are important metrics, they measure capacity and volume (how many requests the system can handle simultaneously) rather than the duration of an individual request. Request size (Option C) measures the payload volume, which can influence latency but does not define the time duration itself. To address the specific requirement of measuring " overall duration, " monitoring request latency within Cloud Monitoring is the standard architectural practice.

NEW QUESTION: 91
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 GKE CPU Utilization and Memory Utilization as service level indicators.
B. Create Request Latency and Error Rate as service level indicators.
C. Create Server Uptime and Error Rate as service level indicators.
D. Create CPU Utilization and Request Latency as service level indicators.
Answer: (SHOW ANSWER)

Valid Professional-Cloud-Architect Dumps shared by ExamDiscuss.com for Helping Passing Professional-Cloud-Architect Exam! ExamDiscuss.com now offer the **newest Professional-Cloud-Architect exam dumps**, the ExamDiscuss.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Professional-Cloud-Architect dumps with Test Engine here: https://www.examdiscuss.com/Google/exam/Professional-Cloud-Architect/premium/ (**353** Q&As Dumps, **35%OFF** **Special Discount Code:** **freecram**)

NEW QUESTION: 92
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)
https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#transfer-options
https://cloud.google.com/storage-transfer-service

NEW QUESTION: 93
Your organization uses Google Kubernetes Engine (GKE) and Amazon Elastic Kubernetes Service (EKS) to manage a complex Kubernetes environment across multiple cloud providers. You need to deploy a solution that streamlines configuration management, enforces security policies, and ensures consistent application deployment across all of the environments. You want to follow Google-recommended practices. What should you do?

A. Leverage Argo CD for GitOps-based continuous delivery and Open Policy Agent (OPA) for policy enforcement, and develop a controller for multi-cluster configuration management.
B. Deploy Customize for configuration customization. Config Sync with multiple Git repositories, and a script to enforce security policies.
C. Utilize Config Sync as part of GKE to synchronize configurations from a centralized repository, and utilize Policy Controller to enforce policies using OPA Gatekeeper.
D. Deploy Crossplane for managing cloud resources as Kubernetes objects. FluxCD for GitOps-based configuration synchronization, and Kyverno for policy enforcement.
Answer: C (LEAVE A REPLY)
According to Google Cloud ' s Anthos (now GKE Enterprise) documentation, Config Sync and Policy Controller are the official tools for multi-cluster and multi-cloud management. These tools allow for a GitOps approach to infrastructure and configuration.
* Config Sync: Automatically synchronizes Kubernetes configurations (like YAML files for Namespaces, Roles, or Deployments) from a central Git repository to all registered clusters, whether they are in GKE or EKS. This ensures consistency across providers.
* Policy Controller: Built on the open-source Gatekeeper project (using Open Policy Agent), it enforces security policies and compliance at scale. For exam ple, it can prevent a developer from deploying a container that runs as root or requires excessive resources across both AWS and Google Cloud environments.
While Option A uses similar open-source technologies, Option C is the " Google-recommended " integrated solution that is fully managed and supported. This centralized management reduces the complexity of maintaining custom controllers (as mentioned in A) or scripts (as mentioned in B), providing a robust, enterprise-grade framework for hybrid and multi-cloud governance.

NEW QUESTION: 94
You are working in a highly secured environment where public Internet access from the Compute Engine VMs is not allowed. You do not yet have a VPN connection to access an on-premises file server. You need to install specific software on a Compute Engine instance. How should you install the software?
A. Upload the required installation files to Cloud Storage. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gsutil.
B. Upload the required installation files to Cloud Storage and use firewall rules to block all traffic except the IP address range for Cloud Storage. Download the files to the VM using gsutil.
C. Upload the required installation files to Cloud Source Repositories. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gcloud.
D. Upload the required installation files to Cloud Source Repositories and use firewall rules to block all traffic except the IP address range for Cloud Source Repositories. Download the files to the VM using gsutil.
Answer: (SHOW ANSWER)
https://cloud.google.com/vpc/docs/private-access-options#pga-supported

NEW QUESTION: 95
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)
https://cloud.google.com/blog/products/databases/getting-started-with-time-series-trend-predictions-using-gcp

NEW QUESTION: 96
Your customer runs a web service used by e-commerce sites to offer product recommendations to users. The company has begun experimenting with a machine learning model on Google Cloud Platform to improve the quality of results.
What should the customer do to improve their model's results over time?
A. Export Cloud Machine Learning Engine performance metrics from Stackdriver to BigQuery, to be used to analyze the efficiency of the model.
B. Build a roadmap to move the machine learning model training from Cloud GPUs to Cloud TPUs, which offer better results.
C. Monitor Compute Engine announcements for availability of newer CPU architectures, and deploy the model to them as soon as they are available for additional performance.
D. Save a history of recommendations and results of the recommendations in BigQuery, to be used as training data.
Answer: (SHOW ANSWER)
https://cloud.google.com/solutions/building-a-serverless-ml-model

NEW QUESTION: 97
You need to reduce the number of unplanned rollbacks of erroneous production deployments in your company's web hosting platform. Improvement to the QA/Test processes accomplished an 80% reduction.
Which additional two approaches can you take to further reduce the rollbacks? Choose 2 answers.
A. Fragment the monolithic platform into microservices
B. Introduce a green-blue deployment model
C. Replace the platform's relational database systems with a NoSQL database
D. Replace the QA environment with canary releases
E. Reduce the platform's dependency on relational database systems
Answer: (SHOW ANSWER)

NEW QUESTION: 98

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

A. Google Compute Engine unmanaged instance groups and Network Load Balancer

B. Google Compute Engine managed instance groups with auto-scaling

C. Google Cloud Dataproc to run Apache Hadoop jobs to process each test

D. Google App Engine with Google Stackdriver for logging

Answer: (SHOW ANSWER)

https://cloud.google.com/compute/d ocs/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: 99

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)

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: 100

For this question, refer to the JencoMart case study
A few days after JencoMart migrates the user credentials database to Google Cloud Platform and shuts down the old server, the new database server stops responding to SSH connections. It is still serving database requests to the application servers correctly. What three steps should you take to diagnose the problem?
Choose 3 answers

Freecram.net

A. Delete the virtual machine (VM) and disks and create a new one.
B. Delete the instance, attach the disk to a new VM, and investigate.
C. Take a snapshot of the disk and connect to a new machine to investigate.
D. Check inbound firewall rules for the network the machine is connected to.
E. Connect the machine to another network with very simple firewall rules and investigate.
F. Print the Serial Console output for the instance for troubleshooting, activate the interactive console, and investigate.
Answer: (SHOW ANSWER)
https://cloud.google.com/compute/docs/troubleshooting/troubleshooting-ssh D: Handling " Unable to connect on port 22 " error message Possible causes include:
There is no firewall rule allowing SSH access on the port. SSH access on port 22 is enabled on all Compute Engine instances by default. If you have disabled access, SSH from the Browser will not work. If you run sshd on a port other than 22, you need to enable the access to that port with a custom firewall rule.
The firewall rule allowing SSH access is enabled, but is not configured to allow connections from GCP Console services. Source IP addresses for browser-based SSH sessions are dynamically allocated by GCP Console and can vary from session to session.
References:
https://cloud.google.com/compute/docs/ssh-in-browser
https://cloud.google.com/compute/docs/ssh-in-browser

NEW QUESTION: 101
Your company is planning to perform a lift and shift migration of their Linux RHEL 6.5+ virtual machines. The virtual machines are running in an on-premises VMware environment. You want to migrate them to Compute Engine following Google-recommended practices. What should you do?
A. 1. Define a migration plan based on the list of the applications and their dependencies.
2. Migrate all virtual machines into Compute Engine individually with Migrate for Compute Engine.
B. 1. Perform an assessment of virtual machines running in the current VMware environment.
2.Create images of all disks. Import disks on Compute Engine.
3.Create standard virtual machines where the boot disks are the ones you have imported.
C. 1. Perform an assessment of virtual machines running in the current VMware environment.
2. Define a migration plan, prepare a Migrate for Compute Engine migration RunBook, and execute the migration.
D. 1. Perform an assessment of virtual machines running in the current VMware environment.
2.Install a third-party agent on all selected virtual machines.
3.Migrate all virtual machines into Compute Engine.
Answer: (SHOW ANSWER)

The framework illustrated in the preceding diagram has four phases:
*Assess. In this phase, you assess your source environment, assess the workloads that you want to migrate to Google Cloud, and assess which VMs support each workload.
*Plan. In this phase, you create the basic infrastructure for Migrate for Compute Engine, such as provisioning the resource hierarchy and setting up network access.
*Deploy. In this phase, you migrate the VMs from the source environment to Compute Engine.
*Optimize. In this phase, you begin to take advantage of the cloud technologies and capabilities.
Reference: https://cloud.google.com/architecture/migrating-vms-migrate-for-compute-engine-getting-started

NEW QUESTION: 102
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

NEW QUESTION: 103
For this question, refer to the Helicopter Racing League (HRL) case study. Your team is in charge of creating a payment card data vault for card numbers used to bill tens of thousands of viewers, merchandise consumers, and season ticket holders. You need to implement a custom card tokenization service that meets the following requirements:
* It must provide low latency at minimal cost.
* It must be able to identify duplicate credit cards and must not store plaintext card numbers.
* It should support annual key rotation.
Which storage approach should you adopt for your tokenization service?
A. Store the card data in Secret Manager after running a query to identify duplicates.
B. Encrypt the card data with a deterministic algorithm stored in Firestore using Datastore mode.
C. Encrypt the card data with a deterministic algorithm and shard it across multiple Memorystore instances.
D. Use column-level encryption to store the data in Cloud SQL.

Answer: (SHOW ANSWER)
https://cloud.google.com/community/tutorials/pci-tokenizer Deterministic output means that a given set of inputs (card number, expiration, and userID) will always generate the same token. This is useful if you want to rely on the token value to deduplicate your token stores. You can simply match a newly generated token to your existing catalog of tokens to determine whether the card has been previously stored. Depending on your application architecture, this can be a very useful feature. However, this could also be accomplished using a salted hash of the input values.

NEW QUESTION: 104
You need to upgrade the EHR connection to comply with their requirements. The new connection design must support business-critical needs and meet the same network and security policy requirements. What should you do?
A. Add a new Dedicated Interconnect connection.
B. Upgrade the bandwidth on the Dedicated Interconnect connection to 100 G.
C. Add a new Carrier Peering connection.
D. Add three new Cloud VPN connections.
Answer: (SHOW ANSWER)

NEW QUESTION: 105
All compute Engine instances in your VPC should be able to connect to an Active Directory server on specific ports. Any other traffic emerging from your instances is not allowed. You want to enforce this using VPC firewall rules.
How should you configure the firewall rules?
A. Create an egress rule with priority 1000 to deny all traffic for all instances. Create another egress rule with priority 100 to allow the Active Directory traffic for all instances.
B. Create an egress rule with priority 100 to deny all traffic for all instances. Create another egress rule with priority 1000 to allow the Active Directory traffic for all instances.
C. Create an egress rule with priority 1000 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 100 to block all traffic for all instances.
D. Create an egress rule with priority 100 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 1000 to block all traffic for all instances.
Answer: (SHOW ANSWER)
https://cloud.google.com/vpc/docs/firewalls

NEW QUESTION: 106
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 Cloud Datastore server to replace the MySQL server in a high-availability configuration. Deploy Jenkins to Compute Engine using Cloud Launcher.

C. Migrate Nginx and Tomcat to App Engine. Deploy a MySQL server using Cloud Launcher. Deploy Jenkins to Compute Engine using Cloud Launcher.
D. Deploy Nginx and Tomcat using Cloud Launcher. Deploy a MySQL server using Cloud Launcher. Deploy Jenkins to Compute Engine using Cloud Deployment Manager scripts.
Answer: (SHOW ANSWER)

Valid Professional-Cloud-Architect Dumps shared by ExamDiscuss.com for Helping Passing Professional-Cloud-Architect Exam! ExamDiscuss.com now offer the **newest Professional-Cloud-Architect exam dumps**, the ExamDiscuss.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Professional-Cloud-Architect dumps with Test Engine here: https://www.examdiscuss.com/Google/exam/Professional-Cloud-Architect/premium/ (**353** Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 107
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: D (LEAVE A REPLY)
https://cloud.google.com/docs /enterprise/best-practices-for-enterprise-organizations#billing _and_management

NEW QUESTION: 108
You are designing a mobile chat application. You want to ensure people cannot spoof chat messages, by providing a message were sent by a specific user.
What should you do
A. Encrypt the message client side using block-based encryption with a shared key.
B. Tag messages client side with the originating user identifier and the destination user.
C. Use public key infrastructure (PKI) to encrypt the message client side using the originating user's private key.
D. Use a trusted certificate authority to enable SSL connectivity between the client application and the server.
Answer: (SHOW ANSWER)

NEW QUESTION: 109
Your company operates nationally and plans to use GCP for multiple batch workloads, including some that are not time-critical. You also need to use GCP services that are HIPAA-certified and manage service costs.
How should you design to meet Google best practices?
A. Provisioning preemptible VMs to reduce cost. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.
B. Provisioning preemptible VMs to reduce cost. Disable and then discontinue use of all GCP and APIs that are not HIPAA-compliant.
C. Provision standard VMs in the same region to reduce cost. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.
D. Provision standard VMs to the same region to reduce cost. Disable and then discontinue use of all GCP services and APIs that are not HIPAA-compliant.
Answer: B (LEAVE A REPLY)
https://cloud.google.com/security/compliance/hipaa/

NEW QUESTION: 110
Your company's user-feedback portal comprises a standard LAMP stack replicated across two zones. It is deployed in the us-central1 region and uses autoscaled managed instance groups on all layers, except the database. Currently, only a small group of select customers have access to the portal. The portal meets a
99.99% availability SLA under these conditions However next quarter, your company will be making the portal available to all users, including unauthenticated users. You need to develop a resiliency testing strategy to ensure the system maintains the SLA once they introduce additional user load. What should you do?
Freecram.net
A. Capture existing users input, and replay captured user load until resource utilization crosses 80%. Also, derive estimated number of users based on existing users usage of the app, and deploy enough resources to handle 200% of expected load.
B. Create synthetic random user input, replay synthetic load until autoscale logic is triggered on at least one layer, and introduce "chaos" to the system by terminating random resources on both zones.
C. Expose the new system to a larger group of users, and increase group ' size each day until autoscale logic is tnggered on all layers. At the same time, terminate random resources on both zones.
D. Capture existing users input, and replay captured user load until autoscale is triggered on all layers. At the same time, terminate all resources in one of the zones.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 111
Your company has decided to build a backup replica of their on-premises user authentication PostgreSQL database on Google Cloud Platform. The database is 4 TB, and large updates are frequent. Replication requires private address space communication. Which networking approach should you use?
A. Google Cloud Dedicated Interconnect
B. Google Cloud VPN connected to the data center network

C. A NAT and TLS translation gateway installed on-premises
D. A Google Compute Engine instance with a VPN server installed connected to the data center network
Answer: (SHOW ANSWER)
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations Google Cloud Dedicated Interconnect provides direct physical connections and RFC 1918 communication between your on-premises network and Google's network. Dedicated Interconnect enables you to transfer large amounts of data between networks, which can be more cost effective than purchasing additional bandwidth over the public Internet or using VPN tunnels.
Benefits:
Traffic between your on-premises network and your VPC network doesn't traverse the public Internet. Traffic traverses a dedicated connection with fewer hops, meaning there are less points of failure where traffic might get dropped or disrupted.
Your VPC network's internal (RFC 1918) IP addresses are directly accessible from your on-premises network. You don't need to use a NAT device or VPN tunnel to reach internal IP addresses. Currently, you can only reach internal IP addresses over a dedicated connection. To reach Google external IP addresses, you must use a separate connection.
You can scale your connection to Google based on your needs. Connection capacity is delivered over one or more 10 Gbps Ethernet connections, with a maximum of eight connections (80 Gbps total per interconnect).
The cost of egress traffic from your VPC network to your on-premises network is reduced. A dedicated connection is generally the least expensive method if you have a high-volume of traffic to and from Google's network.
References: https://cloud.google.com/interconnect/docs/details/dedicated

NEW QUESTION: 112
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 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
B. 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
C. Use an IP range on Google Cloud that does not overlap with the range you use on-premises
D. Use the same IP range on Google Cloud as you use on-premises
Answer: (SHOW ANSWER)

NEW QUESTION: 113
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

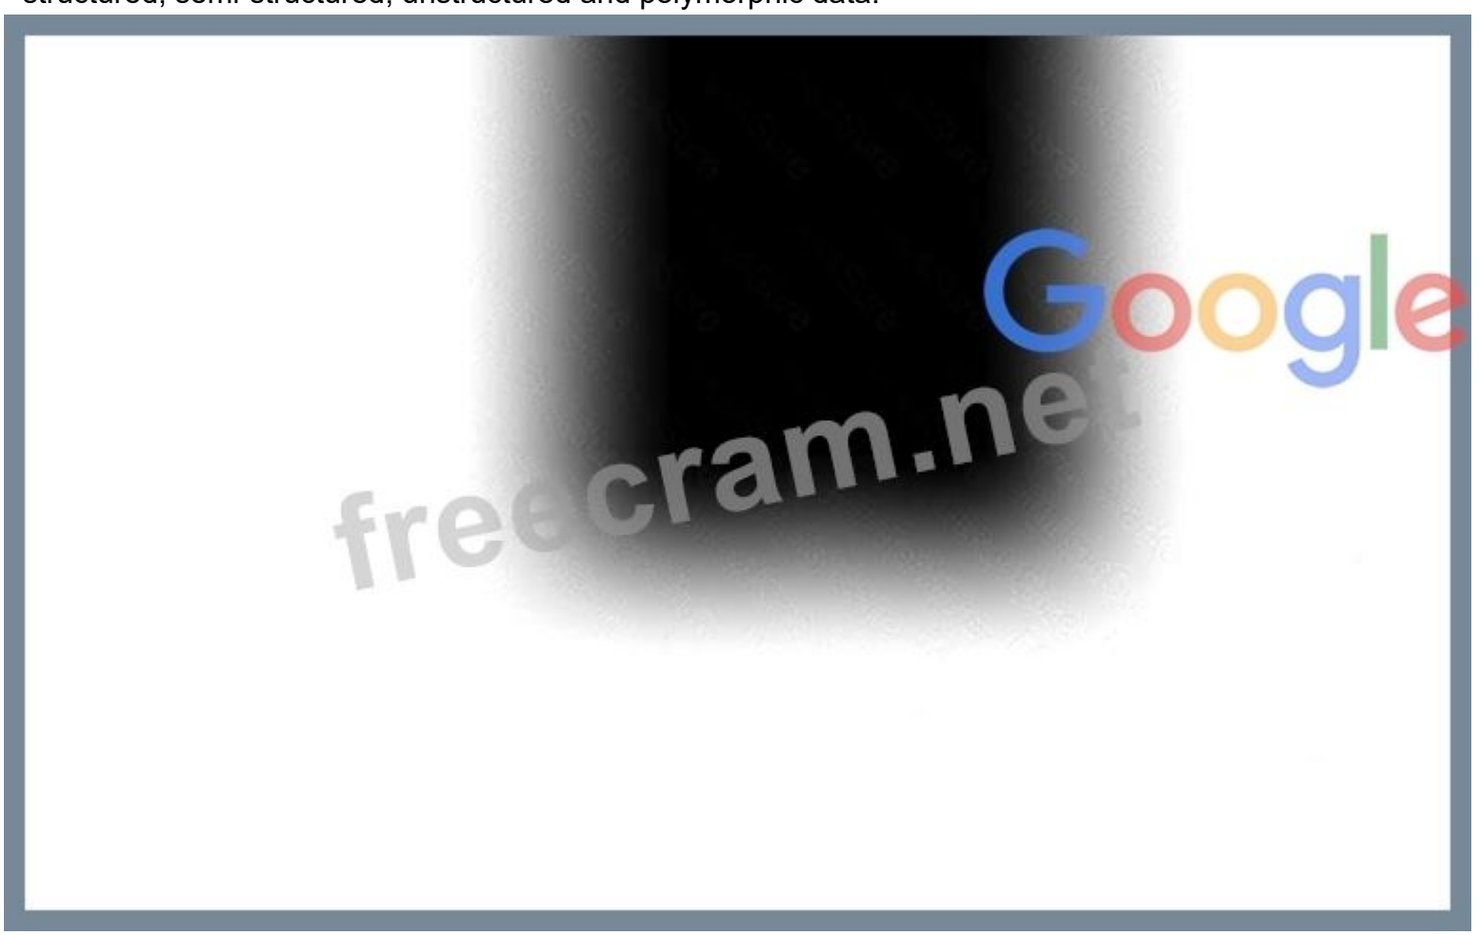
D. Blobstore

Answer: (SHOW ANSWER)

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.



References: https://www.mongodb.com/no sql-explain ed

NEW QUESTION: 114

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: (SHOW ANSWER)
https://cloud.google.com/storage/docs/lifecycle

NEW QUESTION: 115
You are using Cloud SQL as the database backend for a large CRM deployment. You want to scale as usage increases and ensure that you don't run out of storage, maintain 75% CPU usage cores, and keep replication lag below 60 seconds. What are the correct steps to meet your requirements?
A. 1. Enable automatic storage increase for the instance.
2. Create a Stackdriver alert when CPU usage exceeds 75%, and change the instance type to reduce CPU usage.
3. Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
B. 1. Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space.
2. Deploy memcached to reduce CPU load.
3. Change the instance type to a 32-core machine type to reduce replication lag.
C. 1. Enable automatic storage increase for the instance.
2. Change the instance type to a 32-core machine type to keep CPU usage below 75%.
3. Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
D. 1. Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space.
2. Deploy memcached to reduce CPU load.
3. Create a Stackdriver alert for replication lag, and change the instance type to a 32-core machine type to reduce replication lag.
Answer: (SHOW ANSWER)

NEW QUESTION: 116
You have developed a non-critical update to your application that is running in a managed instance group, and have created a new instance template with the update that you want to release. To prevent any possible impact to the application, you don't want to update any running instances. You want any new instances that are created by the managed instance group to contain the new update. What should you do?
A. Start a new rolling restart operation.
B. Start a new rolling replace operation.
C. Start a new rolling update. Select the Proactive update mode.
D. Start a new rolling update. Select the Opportunistic update mode.
Answer: (SHOW ANSWER)
In certain scenarios, an opportunistic update is useful because you don't want to cause instability to the system if it can be avoided. For example, if you have a non-critical update that can be applied as necessary

without any urgency and you have a MIG that is actively being autoscaled, perform an opportunistic update so that Compute Engine does not actively tear down your existing instances to apply the update. When resizing down, the autoscaler preferentially terminates instances with the old template as well as instances that are not yet in a RUNNING state.

NEW QUESTION: 117
For this question, refer to the JencoMart case study.
The JencoMart security team requires that all Google Cloud Platform infrastructure is deployed using a least privilege model with separation of duties for administration between production and development resources.
What Google domain and project structure should you recommend?
A. Create two G Suite accounts to manage users: one for development/test/staging and one for production. Each account should contain one project for every application.
B. Create two G Suite accounts to manage users: one with a single project for all development applications and one with a single project for all production applications.
C. Create a single G Suite account to manage users with each stage of each application in its own project.
D. Create a single G Suite account to manage users with one project for the development/test/staging environment and one project for the production environment.
Answer: (SHOW ANSWER)
Note: The principle of least privilege and separation of duties are concepts that, although semantically different, are intrinsically related from the standpoint of security. The intent behind both is to prevent people from having higher privilege levels than they actually need
* Principle of Least Privilege: Users should only have the least amount of privileges required to perform their job and no more. This reduces authorization exploitation by limiting access to resources such as targets, jobs, or monitoring templates for which they are not authorized.
* Separation of Duties: Beyond limiting user privilege level, you also limit user duties, or the specific jobs they can perform. No user should be given responsibility for more than one related function. This limits the ability of a user to perform a malicious action and then cover up that action.
References: https://cloud.g oogle.com/km s/docs/separation-of-duties

NEW QUESTION: 118
Your company wants to optimize Google Cloud costs for their development and staging environments. These environments are workstations used by developers Monday through Friday, 9:00 AM to 6:00 PM local time.
Currently, the environments run on a fleet of n1-standard-4 Compute Engine instances that operate 24/7, leading to a high monthly cost for resources that are idle more than 70% of the time. You need to implement a solution that significantly reduces the monthly cost of these non-production environments without impacting the development team ' s productivity during work hours. What should you do?
A. Upgrade all instances to the N2 machine series.
B. Purchase three-year committed use discounts (CUDs) for the existing n1-standard-4 instances.

C. Re-architect the environments to run on a regional managed instance group (MIG) with autoscaling enabled.
D. Schedule the virtual machines to start and stop to match your team ' s work schedule.
Answer: (SHOW ANSWER)
The most direct and impactful way to reduce costs for predictable, time-bound workloads is to stop resources when they are not in use . According to the Google Cloud Architecture Framework: Cost Optimization , Compute Engine instances are billed per second while running. By scheduling instances to run only from 9:00 AM to 6:00 PM on weekdays, the instances are active for only 45 hours out of a 168-hour week. This results in a cost reduction of approximately 73% .
While Committed Use Discounts (CUDs) (Option B) provide significant savings (up to 57-70%), they require paying for the resources 24/7 for the duration of the commitment. Even with a 70% discount, paying for 24/7 usage is more expensive than paying full price for only 27% of the week. Autoscaling (Option C) is generally more effective for handling fluctuating traffic loads rather than developer workstations, which are often stateful or individual-specific. Instance Scheduling is a native Compute Engine feature that allows administrators to define start and stop operations automatically, ensuring that productivity is maintained by having the machines ready when developers start their day, while maximizing savings during off-hours.

NEW QUESTION: 119
For this question, refer to the Mountkirk Games case study.
Mountkirk Games' gaming servers are not automatically scaling properly. Last month, they rolled out a new feature, which suddenly became very popular. A record number of users are trying to use the service, but many of them are getting 503 errors and very slow response times. What should they investigate first?
A. Verify that the database is online.
B. Verify that the project quota hasn't been exceeded.
C. Verify that the new feature code did not introduce any performance bugs.
D. Verify that the load-testing team is not running their tool against production.
Answer: (SHOW ANSWER)
503 is service unavailable error. If the database was online everyone would get the 503 error.
https://cloud.google.com/docs/quota#capping_usage

NEW QUESTION: 120
Your company places a high value on being responsive and meeting customer needs quickly. Their primary business objectives are release speed and agility. You want to reduce the chance of security errors being accidentally introduced. Which two actions can you take? Choose 2 answers
Freecram.net
A. Ensure every code check-in is peer reviewed by a security SME.
B. Use source code security analyzers as part of the CI/CD pipeline.
C. Ensure you have stubs to unit test all interfaces between components.
D. Enable code signing and a trusted binary repository integrated with your CI/CD pipeline.
E. Run a vulnerability security scanner as part of your continuous-integration /continuous-delivery (CI /CD) pipeline.
Answer: (SHOW ANSWER)

https://do cs.microsoft.com/en-us/vsts/arti cles/security-valid ation-cicd-pipeline?view=vsts

NEW QUESTION: 121
You have an outage in your Compute Engine managed instance group: all instance keep restarting after 5 seconds. You have a health check configured, but autoscaling is disabled. Your colleague, who is a Linux expert, offered to look into the issue. You need to make sure that he can access the VMs. What should you do?
A. Grant your colleague the IAM role of project Viewer
B. Perform a rolling restart on the instance group
C. Disable the health check for the instance group. Add his SSH key to the project-wide SSH keys
D. Disable autoscaling for the instance group. Add his SSH key to the project-wide SSH Keys
Answer: (SHOW ANSWER)
https://cloud.google.com/compute/docs/instance-groups/autohealing-instances-in-migs Health checks used for autohealing should be conservative so they don't preemptively delete and recreate your instances. When an autohealer health check is too aggressive, the autohealer might mistake busy instances for failed instances and unnecessarily restart them, reducing availability

Valid Professional-Cloud-Architect Dumps shared by ExamDiscuss.com for Helping Passing Professional-Cloud-Architect Exam! ExamDiscuss.com now offer the **newest Professional-Cloud-Architect exam dumps**, the ExamDiscuss.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Professional-Cloud-Architect dumps with Test Engine here: https://www.examdiscuss.com/Google/exam/Professional-Cloud-Architect/premium/ (**353** Q&As Dumps, **35%OFF** **Special Discount Code:** **freecram**)

NEW QUESTION: 122
Refer to the Altostrat Media case study for the following solution regarding the performance analysis of their media processing pipeline.
Altostrat needs to analyze the performance of its media processing pipeline running on Java-based Cloud Run function. You need to select the most effective tool for the task. What should you do?
A. Query logs in Cloud Logging.
B. Analyze the data via Cloud Profiler.
C. Instrument the code to use Cloud Trace.
D. Inspect data from Snapshot Debugger.
Answer: (SHOW ANSWER)
According to Google Cloud documentation on Application Performance Management (APM) , Cloud Profiler is the most effective tool for identifying resource-intensive parts of an application ' s source code.
For Altostrat ' s Java-based media processing pipeline, which likely involves CPU-intensive tasks such as metadata extraction or transcoding, understanding exactly which methods or lines of code consume the most CPU or memory is critical.

Cloud Profiler is a statistical, low-overhead profiler that continuously gathers performance data (CPU usage, heap allocation, etc.) from production applications with less than 0.5% impact on performance. It presents this data in an interactive flame graph , allowing developers to pinpoint bottlenecks in the media processing logic.
While Cloud Trace (Option C) is excellent for monitoring latency across distributed services, it primarily tracks the " outer shell " of the request-showing how long a function took to execute or how long an external API call lasted-but does not provide the deep, code-level insight into internal Java method execution that Profiler does. Cloud Logging (Option A) is useful for troubleshooting specific events but is not a performance analysis tool, and Snapshot Debugger (Option D) is designed for inspecting variables at specific execution points rather than aggregate performance profiling. Given Altostrat ' s goal to " Optimize cloud storage costs " and " Accelerate... operational workflows, " using Cloud Profiler to create more efficient, less resource-hungry processing code is the architecturally sound choice.

NEW QUESTION: 123
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)
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: 124
You are configuring the cloud network architecture for a newly created project m Google Cloud that will host applications in Compote Engine Compute Engine virtual machine instances will be created in two different subnets (sub-a and sub-b) within a single region
* Instances in sub-a win have public IP addresses
* Instances in sub-b will have only private IP addresses
To download updated packages, instances must connect to a public repository outside the boundaries of Google Cloud You need to allow sub-b to access the external repository. What should you do?
A. Enable Private Google Access on sub-b
B. Configure Cloud NAT and select sub b m the NAT mapping section
C. Configure a bastion host instance in sub a to connect to instances in sub-b
D. Enable Identity Aware Proxy for TCP forwarding for instances in sub-b

Answer: (SHOW ANSWER)

* Cloud NAT (network address translation) lets Google Cloud virtual machine (VM) instances without external IP addresses and private Google Kubernetes Engine (GKE) clusters send outbound packets to the internet and receive any corresponding established inbound response packets 1 . By configuring Cloud NAT and selecting sub-b in the NAT mapping section, you can allow instances in sub-b to access the external repository without exposing them to the internet 1 .

NEW QUESTION: 125

You want to establish a Compute Engine application in a single VPC across two regions. The application must communicate over VPN to an on-premises network. How should you deploy the VPN?

A. Use VPC Network Peering between the VPC and the on-premises network.

B. Expose the VPC to the on-premises network using IAM and VPC Sharing.

C. Create a global Cloud VPN Gateway with VPN tunnels from each region to the on-premises peer gateway.

D. Deploy Cloud VPN Gateway in each region. Ensure that each region has at least one VPN tunnel to the on-premises peer gateway.

Answer: (SHOW ANSWER)

https://cloud.google.com/vpn/docs/how-to/creating-static-vpns

NEW QUESTION: 126

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: (SHOW ANSWER)

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: 127

You want to store critical business information in Cloud Storage buckets. The information is regularly changed but previous versions need to be referenced on a regular basis. You want to ensure that there is a record of all changes to any information in these buckets. You want to ensure that accidental edits or deletions can be easily roiled back. Which feature should you enable?

A. Object change notification

B. Bucket Lock

C. Object Lifecycle Management

D. Object Versioning
Answer: (SHOW ANSWER)

NEW QUESTION: 128
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 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.
B. Use Cloud Dataprep and configure the BigQuery tables as the source. Schedule a daily job to clean the data.
C. Set up a streaming Cloud Dataflow job, receiving data by the ingestion process. Clean the data in a Cloud Dataflow pipeline.
D. Create a Cloud Function that reads data from BigQuery and cleans it. Trigger it. Trigger the Cloud Function from a Compute Engine instance.
Answer: (SHOW ANSWER)

NEW QUESTION: 129
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 dat a. You need to design the solution to meet this requirement. What should you do?
A. Create a snapshot schedule for the disk containing the application data. Whenever a zonal outage occurs, use the latest snapshot to restore the disk in the same zone.
B. Configure the Compute Engine instances with an instance template for the application, and use a regional persistent disk for the application data. Whenever a zonal outage occurs, use the instance template to spin up the application in another zone in the same region. Use the regional persistent disk for the application data.
C. Create a snapshot schedule for the disk containing the application data. Whenever a zonal outage occurs, use the latest snapshot to restore the disk in another zone within the same region.
D. Configure the Compute Engine instances with an instance template for the application, and use a regional persistent disk for the application data. Whenever a zonal outage occurs, use the instance template to spin up the application in another region. Use the regional persistent disk for the application data,
Answer: (SHOW ANSWER)
Regional persistent disk is a storage option that provides synchronous replication of data between two zones in a region. Regional persistent disks can be a good building block to use when you implement HA services in Compute Engine. https://cloud.google.com/compute/docs/disks/high-availability-regional-persistent-disk

NEW QUESTION: 130

You are managing an application deployed on Cloud Run for Anthos, and you need to define a strategy for deploying new versions of the application. You want to evaluate the new code with a subset of production traffic to decide whether to proceed with the rollout. What should you do?

Freecram.net

A. Deploy a new revision to Cloud Run with the new version. Configure traffic percentage between revisions.

B. Deploy a new service to Cloud Run with the new version. Add a Cloud Load Balancing instance in front of both services.

C. In the Google Cloud Console page for Cloud Run, set up continuous deployment using Cloud Build for the development branch. As part of the Cloud Build trigger, configure the substitution variable TRAFFIC_PERCENTAGE with the percentage of traffic you want directed to a new version.

D. In the Google Cloud Console, configure Traffic Director with a new Service that points to the new version of the application on Cloud Run. Configure Traffic Director to send a small percentage of traffic to the new version of the application.

Answer: (SHOW ANSWER)

https://cloud.google.com/run/docs/rollouts-rollbacks-traffic-migration

NEW QUESTION: 131

Your organization has a significant amount of log data stored in Cloud Logging. The data engineering team is accustomed to using SQL for analysis and wants the ability to create insightful dashboards for visualizing log trends and patterns. You want to follow the recommendations of the Google Cloud Well-Architected Framework to provide a solution for the data engineering team What should you do?

A. Enable log analytics and run queries in the linked log dataset in BigQuery. Visualize the data with Looker Studio dashboards.

B. Create a log sink, and export the data to a storage bucket. Create an external table in BigQuery for the data in the bucket Run queries and visualize the data with Cloud Monitoring dashboards

C. Enable log analytics and run queries in Cloud Monitoring Visualize the data using Vertex AI workbench.

D. Create a log sink, and export the data to BigQuery using Pub/Sub. Run queries and visualize the data with Cloud Monitoring dashboards.

Answer: (SHOW ANSWER)

The Google Cloud Well-Architected Framework recommends using specialized tools for specialized tasks.

* Log Analytics: This feature allows you to query your Cloud Logging data using Standard SQL directly in the console.

* BigQuery & Looker Studio: By linking the log bucket to BigQuery, your data engineering team can use their existing SQL skills to perform complex joins. Looker Studio (Option A) is then used to create high-level visualizations that are far more advanced than basic Cloud Monitoring charts.

* Why not D? While Pub/Sub to BigQuery works, it is more complex to manage (" operational complexity ") than simply enabling the built-in Log Analytics feature.

NEW QUESTION: 132

You are deploying a critical application with a stateless, containerized (rented on Cloud Run and a Cloud SQL tor PostgreSQL backend The application experiences unpredictable traffic spikes, and the business

requires the ability to immediately roll back a failed deployment to the last known good state You need to apply a deployment strategy that aligns with Site Reliability Engineering (SRE) principles for both the application code and the database schema updates, while meeting the business ' s requirements What should you do?

A. Configure the CI/CD pipeline to use the: latest Last container tag for deployments, with database schema changes applied manually as needed.

B. Use a single CI/CD pipeline that first applies database schema changes and then deploys the new Cloud Run revision

C. Separate CI/CD pipelines for database schema migrations from application deployments When deploying a new Cloud Run revision, use gradual traffic split.

D. Package the database schema migration script within the container to be executed on every container startup before the application process begins.

Answer: (SHOW ANSWER)

In SRE (Site Reliability Engineering), the goal is to decouple components to reduce the " blast radius " of changes.

* Separation of Concerns: Database schema migrations and application code should have separate pipelines (Option C). If a schema update is backward-compatible (a requirement for zero-downtime), the app can be rolled back independently.

* Traffic Management: Cloud Run supports gradual traffic splitting , allowing you to send 5% of traffic to a new revision. If errors spike, you can immediately shift 100% of traffic back to the old revision.

* Why not D? Running migrations on container startup is a common " anti-pattern " in large-scale systems because if 100 containers start at once, they might all try to modify the schema simultaneously, leading to lock contention or database corruption.

NEW QUESTION: 133

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

A. Configure their replication to use UDP.

B. Send the replicated transaction to Google Cloud Pub/Sub.

C. Add additional VPN connections and load balance them.

D. Configure a Google Cloud Dedicated Interconnect.

E. Restore their database daily using Google Cloud SQL.

Answer: (SHOW ANSWER)

NEW QUESTION: 134

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: 135
Your company is expanding its AI-powered operations nationwide and has chosen accelerator-based compute for the AI workloads. The batch image processing workloads are not time-sensitive and can tolerate interruptions. You need to rapidly deploy cost-effective accelerator nodes for these batch tasks, ensuring data persistence when necessary. What should you do?
A. Deploy Cloud Run functions with ephemeral local SSD.
B. Deploy standard VMs with configured accelerators and attached persistent disks.
C. Deploy spot VMs with local SSD to reduce time for burst workloads.
D. Deploy spot VMs with attached persistent disks and implement checkpoint mechanisms.
Answer: (SHOW ANSWER)
For batch processing workloads that are " not time-sensitive " and " can tolerate interruptions, " Spot VMs (formerly Preemptible VMs) are the Google-recommended best practice for cost optimization. Spot VMs offer the same performance as standard VMs (including support for GPUs and TPUs) at a discount of up to 91% .
However, because Spot VMs can be reclaimed by Google Cloud at any time, the architecture must be fault-tolerant . According to Vertex AI and Compute Engine best practices , using Persistent Disks (rather than Local SSDs, which are wiped upon VM termination) ensures that data is not lost when an instance is preempted. By implementing a checkpointing mechanism , the AI model periodically saves its state (weights, processed image metadata, etc.) to the persistent disk. When a new Spot VM is automatically provisioned by a Managed Instance Group to replace the preempted one, the application can resume from the last saved checkpoint rather than starting from the beginning.
Option C is less ideal because Local SSD data does not persist through preemption/reclamation cycles. Option B uses standard VMs, which fails the requirement for " cost-effective " deployment. Option D perfectly balances the extreme cost savings of Spot provisioning with the reliability requirements of large-scale batch image processing.

NEW QUESTION: 136
For this question, refer to the EHR Healthcare case study. EHR has single Dedicated Interconnect connection between their primary data center and Googles network. This connection satisfies EHR's network and security policies:
* On-premises servers without public IP addresses need to connect to cloud resources without public IP addresses
* Traffic flows from production network mgmt. servers to Compute Engine virtual machines should never traverse the public internet.

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

A. Add a new Dedicated Interconnect connection

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

C. Add three new Cloud VPN connections

D. Add a new Carrier Peering connection

Answer: (SHOW ANSWER)

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

Topic 11, Cymbal Retail case study

Company Overview

Cymbal is an online retailer experiencing significant growth. The retailer specializes in a large assortment of products spanning several retail sub-verticals, which makes managing their extensive product catalog a constant challenge.

Solution Concept

Cymbal wants to modernize its operations and enhance the customer experience in three core areas:

* Catalog and Content Enrichment: Cymbal wants to automate and improve the accuracy of their product catalog by utilizing Gen AI to generate product attributes, descriptions, and images from supplier-provided information. This solution will streamline their catalog management, reduce manual effort and errors, and ensure information is consistent across all their sales channels.
* Conversational Commerce with Product Discovery: To enhance customer engagement and drive sales conversion, Cymbal wants to implement a Conversational Commerce solution. This solution will involve integrating AI-powered virtual agents into their website and mobile app to provide customers with a personalized and intuitive shopping experience through natural language conversations. These agents will utilize Google Cloud ' s Discovery AI to process user requests and retrieve the most relevant products based on each customer ' s needs and preferences, creating a more engaging and satisfying shopping journey.
* Technical Stack Modernization: To streamline operations and reduce costs around manual processes, data transfer, error handling and remediation, Cymbal wants to modernize their technical stack with cloud-based infrastructure, secure and efficient data handling, 3rd party integrations, and proactive monitoring and security.

Existing Technical Environment

Cymbal currently relies on the following environment:

* A mix of on-premises and cloud-based systems.
* A variety of databases, including MySQL, Microsoft SQL Server, Redis, and MongoDB, to store and manage its vast product catalog and customer data.
* Kubernetes clusters to run containerized applications.

* Legacy file-based integrations with on-premises systems, including SFTP file transfers, ETL batch processing.
* A custom-built web application which allows customers to browse the product catalog by querying the relational databases for names and categories of products.
* An IVR (Interactive Voice Response) system to handle initial customer calls and route them to the appropriate departments or agents.
* Call center agents who receive transferred calls from the IVR system and manually enter orders into the system when a customer can ' t complete a transaction on their own.
* Various open source tools for monitoring such as Grafana, Nagios, and Elastic.

The current technical environment has encountered significant challenges: manual processes are time-consuming and error-prone, data silos limit a unified view of the customer journey, and integrating new technologies is difficult.

Business Requirements

* Cymbal has outlined these key business requirements for the Gen Al solution:
* Automate Product Catalog Enrichment: Reduce manual effort, minimize errors, and ensure accuracy and consistency across the product catalog.
* Improve Product Discoverability: Enhance search relevance and enable customers to find products more efficiently.
* Increase Customer Engagement: Create a more interactive and personalized
* shopping experience to improve customer satisfaction and potentially reduce product returns.
* Drive Sales Conversion: Provide a more intuitive and helpful shopping experience to improve sales conversion rates and drive revenue growth.
* Reduce costs: Reduce call center staffing costs and data-center hosting costs.

Technical Requirements

To successfully meet Cymbal ' s business needs, the following technical capabilities are essential:

* Attribute Generation: Accurately derive relevant product attributes from various supplier data, including titles, descriptions, and images, ensuring the attributes align with the product category and Cymbal ' s existing catalog structure.
* Image Generation and Enhancement: Generate different product image variations from a base image (e.g., showcasing various colors). It should also support background changes, product color adjustments, and the addition of text overlays.
* Automate Product Discovery: Process customer requests expressed in natural language and return highly relevant product results.
* Scalability and Performance: The solution must handle Cymbal ' s extensive product catalog and accommodate their anticipated growth without compromising performance or user experience.
* Human-in-the-Loop (HITL) Review: Provide a user interface (UI) for associates to review and manage Gen Al-generated content, allowing them to approve, reject, or modify suggestions before updating the product catalog.
* Data Security and Compliance: Ensure all customer data, including product information and interactions with virtual agents, are handled securely and comply with relevant industry regulations.

Executive Statement

By implementing Google Cloud ' s Gen AI for Digital Commerce solutions, Cymbal can transform its online retail operations to improve efficiency, enhance customer experience, and drive revenue growth. Key benefits for Cymbal include:
* Reduced operational costs through automation of catalog management tasks.
* Increased efficiency and speed in onboarding new products and updating existing ones.
* Improved accuracy and consistency of product information across all sales channels.
* A more engaging and personalized shopping experience that caters to modern
* customer preferences for conversational commerce.
* Enhanced product discoverability leading to higher conversion rates and increased
* sales.
This strategic investment in generative AI will position Cymbal to remain competitive and thrive in the rapidly evolving landscape of online retail.

Valid Professional-Cloud-Architect Dumps shared by ExamDiscuss.com for Helping Passing Professional-Cloud-Architect Exam! ExamDiscuss.com now offer the **newest Professional-Cloud-Architect exam dumps**, the ExamDiscuss.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Professional-Cloud-Architect dumps with Test Engine here: https://www.examdiscuss.com/Google/exam/Professional-Cloud-Architect/premium/ (**353** Q&As Dumps, **35%OFF** **Special Discount Code:** **freecram**)

NEW QUESTION: 137
A rapidly growing online gaming company wants to establish a secure and reliable connection between their on-premises data center and Google Cloud VPC network. They want to achieve a 99.99% availability SLA and keep initial costs down. They need you to quickly implement a solution to support an upcoming game launch, and you must ensure that all traffic in transit is encrypted. What should you do?
A. Implement a Classic VPN connection.
B. Implement an HA VPN connection.
C. Configure a Partner Interconnect connection.
D. Provision a Dedicated Interconnect connection.
Answer: (SHOW ANSWER)
Google Cloud HA (High Availability) VPN is a cloud-native, managed IPsec VPN service designed to deliver a 99.99% network availability SLA. Unlike legacy Classic VPN, HA VPN utilizes an architecture that deploys two external IP addresses assigned to two independent interfaces on a single gateway instance, operating across multiple underlying availability zones. This architecture ensures complete hardware and path redundancy.
Interconnect solutions (Dedicated and Partner) provide dedicated physical layer paths that require extensive provisioning timelines, physical hardware cross-connect setups, and higher upfront capital investments, making them incompatible with the need for rapid deployment before a pending launch. Furthermore, Interconnect traffic is unencrypted by default unless layered with Cloud VPN over

Interconnect. HA VPN satisfies the encryption-in-transit mandate via native IPsec suites, avoids high initial costs by using the public internet as transport, and installs via automated configurations.
References: Cloud Architect (GCP) Study Guide topics: " Cloud VPN Topologies " , " HA VPN Architecture and SLAs " , " IPsec Encryption Standards " , and " Rapid Deployment Hybrid Networking " .

NEW QUESTION: 138
Your company is running a stateless application on a Compute Engine instance. The application is used heavily during regular business hours and lightly outside of business hours. Users are reporting that the application is slow during peak hours. You need to optimize the application's performance. What should you do?
A. Create a snapshot of the existing disk. Create an instance template from the snapshot. Create an autoscaled managed instance group from the instance template.
B. Create a snapshot of the existing disk. Create a custom image from the snapshot. Create an autoscaled managed instance group from the custom image.
C. Create a custom image from the existing disk. Create an instance template from the custom image. Create an autoscaled managed instance group from the instance template.
D. Create an instance template from the existing disk. Create a custom image from the instance template.Create an autoscaled managed instance group from the custom image.
Answer: (SHOW ANSWER)
https://cloud.google.com/compute/docs/instance-templates/create-instance-templates

NEW QUESTION: 139
For this question, refer to the Dress4Win case study.
You want to ensure Dress4Win's sales and tax records remain available for infrequent viewing by auditors for at least 10 years. Cost optimization is your top priority. Which cloud services should you choose?
A. Google Cloud Storage Coldline to store the data, and gsutil to access the data.
B. Google Cloud Storage Nearline to store the data, and gsutil to access the data.
C. Google Bigtabte with US or EU as location to store the data, and gcloud to access the data.
D. BigQuery to store the data, and a web server cluster in a managed instance group to access the data.
Google Cloud SQL mirrored across two distinct regions to store the data, and a Redis cluster in a managed instance group to access the data.
Answer: (SHOW ANSWER)
References: https://cloud.google.com/storage/docs/storage-classes

NEW QUESTION: 140
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?
Freecram.net
A. All admin and VM system logs are automatically collected by Stackdriver.
B. Stackdriver automatically collects admin activity logs for most services. The Stackdriver Logging agent must be installed on each instance to collect system logs.

C. Launch a custom syslogd compute instance and configure your GCP project and VMs to forward all logs to it.
D. Install the Stackdriver Logging agent on a single compute instance and let it collect all audit and access logs for your environment.
Answer: (SHOW ANSWER)
https://cloud.google.com/logging/docs/agent/default-logs

NEW QUESTION: 141
Your product learn is building a critical, customer-facing application on Google Cloud The development team wants to use Spanner for their database to take advantage of its horizontal scalability and low operational overhead However, the FinOps team is concerned about the direct monthly cost of Spanner and proposed using a self-managed PostgreSQL database on Compute Engine VMs instead. You need to resolve this conflict and ensure the project moves forward with an architecturally sound database choice that balances technical requirements with financial constraints. What should you do?
A. Suggest using Cloud SQL for PostgreSQL as a compromise to get a managed service at a lower cost than Spanner.
B. Develop a total cost of ownership (TCO) analysis that includes operational overhead, and present it m a workshop to facilitate a decision.
C. Provide the development team with a reference architecture for deploying a highly available PostgreSQL cluster on a regional managed instance group (MIG).
D. Cite the reliability and performance optimization pillars of the Google Cloud Well-Architected Framework to formally justify the use of Spanner.
Answer: (SHOW ANSWER)
This question highlights the classic conflict between Sticker Price (direct cost) and Operational Expense (OpEx). While a Compute Engine VM running PostgreSQL (Option C) looks cheaper on a monthly bill, it ignores the cost of human labor. A Total Cost of Ownership (TCO) analysis (Option B) is the professional architectural approach to resolve this. It accounts for the salary of DBAs required to manage sharding, patching, and manual failovers-tasks that Spanner handles automatically. For a " critical, customer-facing application " requiring horizontal scale, the TCO will almost always show that Spanner is cheaper in the long run because it reduces the " 900-hour time sink " of manual database maintenance.

NEW QUESTION: 142
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 fewer changes to production.
B. Increase the load on your test and staging environments.
C. Deploy smaller changes to production.
D. Deploy changes to a small subset of users before rolling out to production.
Answer: (SHOW ANSWER)

NEW QUESTION: 143
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)
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: 144
You have an application deployed on Kubernetes Engine using a Deployment named echo-deployment. The deployment is exposed using a Service called echo-service. You need to perform an update to the application with minimal downtime to the application. What should you do?
A. Use kubect1 set image deployment/echo-deployment <new-image>
B. Use the rolling update functionality of the Instance Group behind the Kubernetes cluster
C. Update the deployment yaml file with the new container image. Use kubect1 delete deployment/ echo-deployment and kubect1 create -f <yaml-file>
D. Update the service yaml file which the new container image. Use kubect1 delete service/echoservice and kubect1 create -f <yaml-file>
Answer: (SHOW ANSWER)
https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps#updating_an_application

NEW QUESTION: 145
Your company wants to migrate their 10-TB on-premises database export into Cloud Storage You want to minimize the time it takes to complete this activity, the overall cost and database load. The bandwidth

between the on-premises environment and Google Cloud is 1 Gbps. You want to follow Google-recommended practices What should you do?

A. Use the Data Transfer appliance to perform an offline migration

B. Use a commercial partner ETL solution to extract the data from the on-premises database and upload it into Cloud Storage

C. Develop a Dataflow job to read data directly from the database and write it into Cloud Storage

D. Compress the data and upload it with gsutii -m to enable multi-threaded copy

Answer: (SHOW ANSWER)

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

* https://cloud.google.com/data-transfer-appliance/docs/overview

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

NEW QUESTION: 146

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)

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: 147
A lead engineer wrote a custom tool that deploys virtual machines in the legacy data center. He wants to migrate the custom tool to the new cloud environment You want to advocate for the adoption of Google Cloud Deployment Manager What are two business risks of migrating to Cloud Deployment Manager? Choose 2 answers
A. Cloud Deployment Manager uses Python.
B. Cloud Deployment Manager APIs could be deprecated in the future.
C. Cloud Deployment Manager is unfamiliar to the company's engineers.
D. Cloud Deployment Manager requires a Google APIs service account to run.
E. Cloud Deployment Manager can be used to permanently delete cloud resources.
F. Cloud Deployment Manager only supports automation of Google Cloud resources.
Answer: (SHOW ANSWER)
https://cloud.google.com/deployment-manager/docs/deployments/deleting-deployments

NEW QUESTION: 148
You have found an error in your App Engine application caused by missing Cloud Datastore indexes. You have created a YAML file with the required indexes and want to deploy these new indexes to Cloud Datastore.
What should you do?
A. Point gcloud datastore create-indexes to your configuration file
B. Upload the configuration file the App Engine's default Cloud Storage bucket, and have App Engine detect the new indexes
C. Create an HTTP request to the built-in python module to send the index configuration file to your application
D. In the GCP Console, use Datastore Admin to delete the current indexes and upload the new configuration file
Answer: (SHOW ANSWER)

NEW QUESTION: 149
Mountkirk Games wants you to secure the connectivity from the new gaming application platform to Google Cloud. You want to streamline the process and follow Google-recommended practices. What should you do?
A. Use Kubernetes Secrets, which are obfuscated by default. Configure these Secrets to be used by the application platform.
B. Configure HashiCorp Vault on Compute Engine, and use customer managed encryption keys and Cloud Key Management Service (Cloud KMS) to manage the encryption keys. Configure these Secrets to be used by the application platform.

C. Configure Kubernetes Secrets to store the secret, enable Application-Layer Secrets Encryption, and use Cloud Key Management Service (Cloud KMS) to manage the encryption keys. Configure these Secrets to be used by the application platform.
D. Configure Workload Identity and service accounts to be used by the application platform.
Answer: (SHOW ANSWER)

NEW QUESTION: 150
Your company is developing a web-based application. You need to make sure that production deployments are linked to source code commits and are fully auditable. What should you do?
Freecram.net
A. Make sure a developer is tagging the code commit with the date and time of commit
B. Make sure a developer is adding a comment to the commit that links to the deployment.
C. Make the container tag match the source code commit hash.
D. Make sure the developer is tagging the commits with :latest
Answer: (SHOW ANSWER)
From: https://cloud.google.com/architecture/best-practices-for-building-containers Under: Tagging using the Git commit hash (bottom of page almost)
"In this case, a common way of handling version numbers is to use the Git commit SHA-1 hash (or a short version of it) as the version number. By design, the Git commit hash is immutable and references a specific version of your software.
You can use this commit hash as a version number for your software, but also as a tag for the Docker image built from this specific version of your software. Doing so makes Docker images traceable: because in this case the image tag is immutable, you instantly know which specific version of your software is running inside a given container."

Valid Professional-Cloud-Architect Dumps shared by ExamDiscuss.com for Helping Passing Professional-Cloud-Architect Exam! ExamDiscuss.com now offer the **newest Professional-Cloud-Architect exam dumps**, the ExamDiscuss.com Professional-Cloud-Architect exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Professional-Cloud-Architect dumps with Test Engine here: https://www.examdiscuss.com/Google/exam/Professional-Cloud-Architect/premium/ (**353** Q&As Dumps, **35%OFF** **Special Discount Code:** **freecram**)