Google

Exam Questions Cloud-Digital-Leader

Google Cloud Digital Leader exam

NEW QUESTION 1
- (Topic 1)
Your manager wants to restrict communication of all virtual machines with internet access; with resources in another network; or with a resource outside Compute Engine. It is expected that different teams will create new folders and projects in the near future.
How would you restrict all virtual machines from having an external IP address?

A. Define an organization policy at the root organization node to restrict virtual machine instances from having an external IP address
B. Define an organization policy on all existing folders to define a constraint to restrict virtual machine instances from having an external IP address
C. Define an organization policy on all existing projects to restrict virtual machine instances from having an external IP address
D. Communicate with the different teams and agree that each time a virtual machine is created, it must be configured without an external IP address

Answer: A

Explanation:
Reference: https://cloud.google.com/resource-manager/docs/organization-policy/overview

NEW QUESTION 2
- (Topic 1)
An organization has completely migrated all their infrastructure to the cloud to benefit from its agility. Now they want to innovate faster and achieve a higher return on investment. What should the organization do?

A. Manually provision all cloud infrastructure for increased control.
B. Modernize their applications.
C. Lower their service level objective (SLO).
D. Move to a hybrid architecture with some of their infrastructure on-premises.

Answer: B

Explanation:
Because this will enable the business to better serve their users.

NEW QUESTION 3
- (Topic 1)
You want to build an application that will allow customers to register and login. It would be great to have the ability to secure it with multi-factor authentication and the ability to reset credentials. As a small startup, you want to build the main application as quickly as possible and have minimum overhead. Which might be a suitable option for you on Google Cloud?

A. Since identity and credentials should be secure and private, do not trust other ser-vice providers.
B. Cloud Identity
C. Google Workspace
D. Cloud Identity Platform

Answer: D

Explanation:

Cloud Identity Platform
Cloud Identity Platform allows you to manage identity and credentials for your consumer- facing applications. So that's the right one in this case to use. "Identity Platform is a customer identity and access management (CIAM) platform that helps organizations add identity and access management functionality to their applications, protect user accounts, and scale with confidence on Google Cloud."
Reference link- https://cloud.google.com/identity-platform

NEW QUESTION 4
- (Topic 1)
Your application has repeated data requests of the exact same nature. At the same time, the number of user requests is increasing. Monitoring indicates that the load on the existing database is increasing, and there seems to be a bottleneck. An analysis of the data requested shows us that it is application-managed data and that it changes, but not often. How can you improve the efficiency of the application?

A. Use Cloud Memorystore to improve speed via caching
B. Increase the amount of RAM on the machine hosting the database so that it has higher data throughput.
C. Use Cloud Storage with multi-regional storage so that all users accessing the data will have lower latency
D. Increase the number of CPUs on the machine hosting the database so that it has higher data throughput.

Answer: A

Explanation:
Cloud Memorystore is an in-memory database that has sub-millisecond latency. This is ideal for caching application data that also changes once in a while.
https://cloud.google.com/memorystore

NEW QUESTION 5
- (Topic 1)
Your organization is building an application running in Google Cloud. Currently, software builds, tests, and regular deployments are done manually, but you want to reduce work for the team. Your organization wants to use Google Cloud managed solutions to automate your build, testing, and deployment process.
Which Google Cloud product or feature should your organization use?

A. Cloud Scheduler
B. Cloud Code

C. Cloud Build
D. Cloud Deployment Manager

Answer: C

Explanation:
Deploy your application to App Engine using the gcloud app deploy command. This command automatically builds a container image by using the Cloud Build service and then deploys that image to the App Engine flexible environment.
Reference: https://cloud.google.com/appengine/docs/flexible/nodejs/testing-and-deploying-your-app

NEW QUESTION 6
- (Topic 1)
Your organization needs to establish private network connectivity between its on-premises network and its workloads running in Google Cloud. You need to be able to set up the connection as soon as possible.
Which Google Cloud product or feature should you use?

A. Cloud Interconnect
B. Direct Peering
C. Cloud VPN
D. Cloud CDN

Answer: A

Explanation:
Private Google Access for on-premises hosts provides a way for on-premises systems to connect to Google APIs and services by routing traffic through a Cloud VPN tunnel. Reference: https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid

NEW QUESTION 7
- (Topic 1)
Each of the three cloud service models - infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS) - offers benefits between flexibility and levels of management by the cloud provider and the customer.
Why would SaaS be the right choice of service model?

A. You want a balance between flexibility for the customer and the level of management by the cloud provider
B. You want to minimize the level of management by the customer
C. You want to maximize flexibility for the customer.
D. You want to be able to shift your emphasis between flexibility and management by the cloud provider as business needs change

Answer: B

Explanation:
Benefits of SaaS
The main benefit of SaaS is that it offloads all infrastructure and application management to the SaaS vendor.
Reference: https://www.ibm.com/cloud/learn/iaas-paas-saas

NEW QUESTION 8
- (Topic 1)
An IoT platform is providing services to home security systems. They have more than a million customers, each with many home devices. Burglaries or child safety issues are concerns that the clients customers. Therefore, the platform has to respond very quickly in near real time. What could be a typical data pipeline used to support this platform on Google Cloud?

A. Cloud Pub/Sub, Cloud Dataflow, Data Studio
B. Cloud Functions, Cloud Dataproc, Looker
C. Cloud Pub/Sub, Cloud Dataflow, BigQuery
D. Cloud Functions, Cloud Dataproc, BigQuery

Answer: A

Explanation:
Explanation
=> Cloud Pub/Sub- Cloud Pub/Sub is the best to be the end-point for ingesting large amounts of data. It will grow as required, can stream data to downstream systems, and can also work with intermittently available backends.
=> Cloud Dataflow- supports streaming data and therefore is an appropriate option for processing the data that is ingested.
=> BigQuery- BigQuery also supports streaming data and its possible to do real time ana- lytics on it.
=> DataStudio- DataStudio and Looker are for visualization. They don't have any in-built analysis.
=> Cloud Functions- Cloud Functions is a useful serverless endpoint. However, Pub/Sub is better in this case because it can also retain messages for a set period if it was not possi- ble to deliver it first time.
=>Cloud Dataproc- Cloud Dataproc is used for Hadoop/Spark workloads and won't be a good fit here.

NEW QUESTION 9
- (Topic 1)
Your organization is running all its workloads in a private cloud on top of a hypervisor. Your
organization has decided it wants to move to Google Cloud as quickly as possible. Your organization wants minimal changes to the current environment, while using the maximim amount of managed services Google offers.
What should your organization do?

A. Migrate the workloads to Google Cloud VMware Engine
B. Migrate the workloads to Compute Engine
C. Migrate the workloads to Bare Metal Solution

D. Migrate the workloads to Google Kubernetes Engine

Answer: B

Explanation:
Migrate for Compute Engine enables you to lift and shift workloads at scale to Google Cloud Compute Engine with minimal changes and risk.
Reference: https://dataintegration.info/simplify-vm-migrations-with-migrate-for-compute-engine-as-a-service

NEW QUESTION 10
- (Topic 1)
A startup is planning to create their entire suite of applications on Google Cloud. They are looking at various open source technologies to build applications. One of the considera-tion is about having a well integrated monitoring tool. They have to be able to constantly review load capacity and performance of their applications and virtual machines. What would you advise them to do?

A. It is best to build a custom solution so that they know it integrates well with all their custom applications.
B. Since they are using open source for applications, find another open source monitoring tool and integrate it, which could turn out to be very cheap.
C. Use the Google Cloud Operations Suite which contains monitoring among other operations tools.
D. Update the application code to regularly write to output log
E. Export the logs to BigQuery to analyze them frequently.

Answer: C

Explanation:

Operations Suite is well integrated into Google and it s the recommended option. References: https://cloud.google.com/products/operations

NEW QUESTION 10
- (Topic 1)
Your organization is developing an application that will capture a large amount of data from millions of different sensor devices spread all around the world. Your organization needs a database that is suitable for worldwide, high-speed data storage of a large amount of unstructured data.
Which Google Cloud product should your organization choose?

A. Firestore
B. Cloud Data Fusion
C. Cloud SQL
D. Cloud Bigtable

Answer: D

Explanation:

Reference: https://cloud.google.com/bigtable
Cloud Bigtable is a sparsely populated table that can scale to billions of rows and thousands of columns, enabling you to store terabytes or even petabytes of data. A single value in each row is indexed; this value is known as the row key. Bigtable is ideal for storing very large amounts of single-keyed data with very low latency. It supports high read and write throughput at low latency, and it is an ideal data source for MapReduce operations.
Bigtable is exposed to applications through multiple client libraries, including a supported extension to the Apache HBase library for Java. As a result, it integrates with the existing Apache ecosystem of open-source Big Data software.
Bigtable's powerful back-end servers offer several key advantages over a self-managed HBase installation:
Incredible scalability. Bigtable scales in direct proportion to the number of machines in
your cluster. A self-managed HBase installation has a design bottleneck that limits the performance after a certain threshold is reached. Bigtable does not have this bottleneck, so you can scale your cluster up to handle more reads and writes.
Simple administration. Bigtable handles upgrades and restarts transparently, and it automatically maintains high data durability. To replicate your data, simply add a second cluster to your instance, and replication starts automatically. No more managing replicas or regions; just design your table schemas, and Bigtable will handle the rest for you.
Cluster resizing without downtime. You can increase the size of a Bigtable cluster for a few hours to handle a large load, then reduce the cluster's size again—all without any downtime. After you change a cluster's size, it typically takes just a few minutes under load for Bigtable to balance performance across all of the nodes in your cluster.
Graphical user interface, text, application, email Description automatically generated

NEW QUESTION 12
- (Topic 1)
You decide to migrate your on-premises environment to the cloud. You need to determine which resource components still need to be assigned ownership.
Which two functions are owned by a public cloud provider? (Choose two.)

A. Hardware maintenance
B. Infrastructure architecture
C. Infrastructure deployment automation
D. Hardware capacity management
E. Fixing application security issues

Answer: AD

Explanation:
In a shared responsible model, hardware maintence and capacity management cloud provider is the responsible part.

NEW QUESTION 17
- (Topic 1)
Your customer has reliable information to indicate that they will use a certain amount of computing and analytics. The workloads are critical and they don't want to take a chance with VMs or BigQuery slots being unavailable during a peak period. How can they ensure that they allocate the capacity?

A. Send in the filled form to Google Cloud support to reserve the Compute Engine and BigQuery resources required.
B. Create reservations on Compute Engine and BigQuery.
C. On the day the capacity is required, set a scheduled job that will provision as many resources as required and lock it in.
D. Google Cloud is elastic for resource
E. You cannot reserve resources in advance; it is pay per use.

Answer: B

Explanation:
Create reservations on Compute Engine and BigQuery. You can reserve capacity in advance and use it over a period of time. You could also get a cost advantage.
=> There is no need for involved support. It is self-serve via the console.
=> You can reserve resources in advance when you have the need for it. And when you want to take a pay-per-use approach, that is also possible.
=> It is not a good idea to be lock in/hoard resources; you'll pay unnecessarily for resources. Also, it is difficult to time exactly when the demand will be.
References:
https://cloud.google.com/compute/docs/instances/reserving-zonal-resources https://cloud.google.com/bigquery/docs/reservations-intro

NEW QUESTION 20
- (Topic 1)
Which of the following is/are true about Bare Metal Solutions?

A. Enterprise-grade deployment platform
B. All your existing investment in tooling and best practices will work as is
C. Continue to run any version, and feature set, any database option, and any cus- tomizations (patchsets)
D. All of the Above.

Answer: D

Explanation:
Bare Metal Solution for Oracle
Bring your Oracle workloads to Google Cloud with Bare Metal Solution and jumpstart your cloud journey with minimal risk.
- Continue to run any version, any feature set, any database option, and any customizations (patchsets)
- Enterprise-grade deployment platform
- High availability with Oracle RAC
- Works with any application, any Oracle versions
- All your existing investment in tooling and best practices will work as is

NEW QUESTION 23
- (Topic 1)
As your organization increases its release velocity, the VM-based application upgrades take a long time to perform rolling updates due to OS boot times. You need to make the application deployments faster.
What should your organization do?

A. Migrate your VMs to the cloud, and add more resources to them
B. Convert your applications into containers
C. Increase the resources of your VMs
D. Automate your upgrade rollouts

Answer: B

NEW QUESTION 28
- (Topic 1)
A company with its own private data center has called you in for help with their disaster recovery planning. News of multiple ransomware attacks has made them very anxious. They want to make they are well prepared for such an eventuality. Which of these would be good recommendations?

A. It is better to have redundancy; so, set up another private data center nearby so that you can quickly go over in case of an emergency.
B. It is better to have redundancy; use one or many of the Google Cloud datacenters as a backup location.
C. The one data center is enough, as long as the data is encrypted; attackers won't be able to read the data.
D. The one data center is enough as long as you regularly back up data and save it in another place in the same DC.

Answer: B

Explanation:

A single data center is vulnerable. So any option involving that is not good. Reference Link:- https://www.coresite.com/blog/data-center-redundancy

NEW QUESTION 29
- (Topic 1)
Your organization recently migrated its compute workloads to Google Cloud. You want these workloads in Google Cloud to privately and securely access your large volume of on- premises data, and you also want to minimize latency.
What should your organization do?

A. Use Storage Transfer Service to securely make your data available to Google Cloud
B. Create a VPC between your on-premises data center and your Google resources
C. Peer your on-premises data center to Google’s Edge Network
D. Use Transfer Appliance to securely make your data available to Google Cloud

Answer: C

Explanation:
Graphical user interface, text, application, Word, email

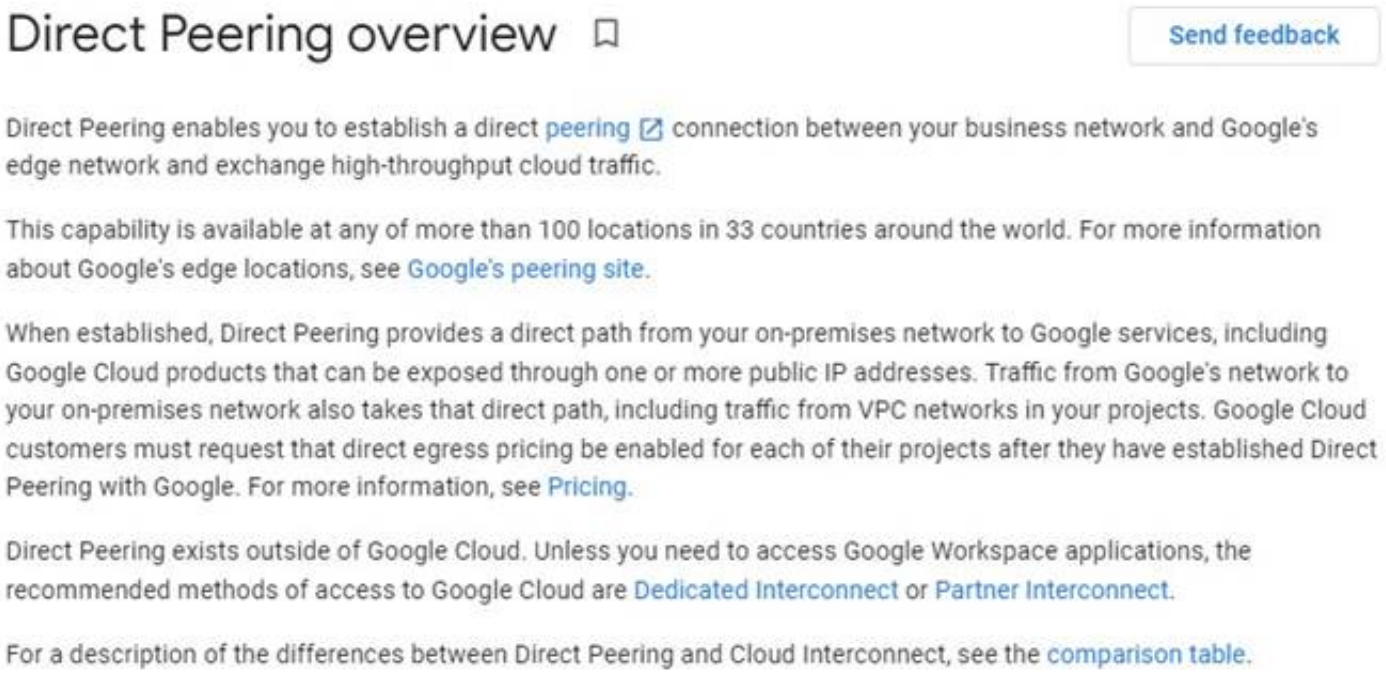

Description automatically generated https://cloud.google.com/network-connectivity/docs/direct-peering

NEW QUESTION 33
- (Topic 1)
Your organization wants to be sure that is expenditures on cloud services are in line with the budget. Which two Google Cloud cost management features help your organization gain greater visibility into its cloud resource costs? (Choose two.)

A. Billing dashboards
B. Resource labels
C. Sustained use discounts
D. Financial governance policies
E. Payments profile

Answer: AB

Explanation:

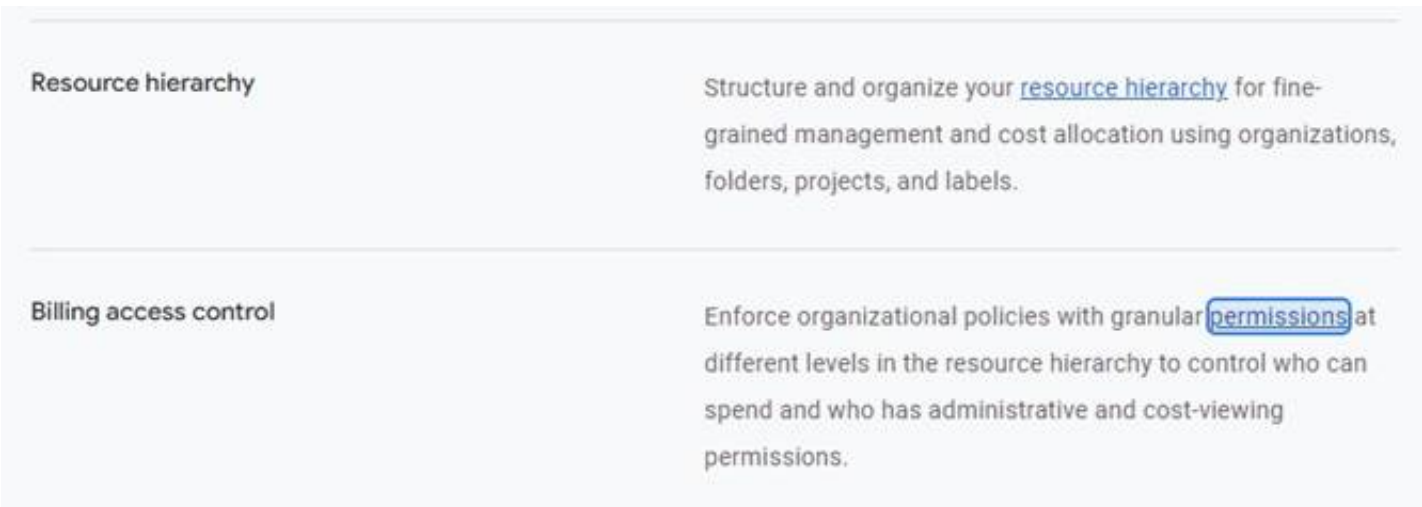

Description automatically generated with medium confidence
A label is a key-value pair that helps you organize your Google Cloud resources. You can attach a label to each resource, then filter the resources based on their labels. Information about labels is forwarded to the billing system, so you can break down your billed charges by label.
Reference link- https://cloud.google.com/cost-management

NEW QUESTION 37
- (Topic 3)
How would an organization benefit from using Looker?

A. Optimal identity and access management
B. Leading serverless warehousing technology
C. Robust data roll-back accuracy
D. Advanced business intelligence and analytics

Answer: D

Explanation:
Looker is a business intelligence software and big data analytics platform that helps you explore, analyze and share real-time business analytics easily.

NEW QUESTION 40
- (Topic 3)
A large retail organization uses traditional technology for their ecommerce website During peaks m traffic resources are often underutilized or overprovisioned
They have decided to migrate to cloud technology
What aspect of cloud technology will benefit their ecommerce business?

A. Agile infrastructure means that they only pay tor what they need, when they need it
B. Shared responsibility means that the cloud provider brings increased visibility during peaks in traffic
C. Operational expenditure means that their total cost of ownership is more predictable
D. Unlimited storage means that their website will never experience downtime

Answer: A

NEW QUESTION 43
- (Topic 3)
An organization is making a strategic change to customer support in response to feedback. They plan to extend their helpline availability hours.
Why is the organization making this change?

A. Users expect professional expertise
B. Users require personalization
C. Users expect always-on services
D. Users require regional access

Answer: C

NEW QUESTION 46
- (Topic 3)
How can a streaming service meet global compliance requirements using the cloud?

A. By automatically encrypting personally identifiable information
B. By obtaining a business license to operate in a new market
C. By allowing users to disable two-factor authentication
D. By adhering only to data polices of the country in which the head office is registered

Answer: A

NEW QUESTION 49
- (Topic 3)
Why is data stored in Google Cloud secure and private?

A. Data is encrypted by the Security Command Center
B. Data is encrypted by Cloud Data Loss Prevention
C. Data is encrypted by default
D. Data is encrypted when an appropriate tag is applied

Answer: C

Explanation:
https://cloud.google.com/docs/security/encryption/default-encryption#:~:text=Google%20uses%20the%20Advanced%20Encryption,to%202015%20that%20use%20AES128

NEW QUESTION 53
- (Topic 3)
After rolling out a new update, an organization found a minor bug in its online video game. How should the organization approach this bug while following SRE principles?

A. Accept and learn from the bug because failure is normal
B. Accept and ignore the bug because it is only minor
C. Hold a postmortem to reprimand the employee responsible for the bug
D. Document bug correction to eliminate all future bugs

Answer: A

Explanation:
https://www.blameless.com/sre/sre-principles
Accepting failure as normal is one of the SRE principles. SREs believe that accepting
failure as normal helps to build an iterative, collaborative culture. One way this is done is by holding a blameless “lessons learned” discussion after an incident occurs.

NEW QUESTION 54
- (Topic 3)
An organization provides a loyalty program for its customers. It recently partnered with other businesses so that customers can get loyalty points at a range of other stores.
Why should the organization use application programming interfaces (APIs)?

A. To migrate all partner data for disaster recovery
B. To analyze and publish loyalty program statistics to a dashboard
C. To personalize recommendations for loyalty card users
D. To connect third-party systems to ensure up-to-date information

Answer: D

NEW QUESTION 58

- (Topic 3)
How does Google Cloud ensure that customer data remains secure and private when at rest?

A. By aggregating training data for customers within each industry
B. By automatically locking files containing suspicious code
C. By auditing platform privacy practices against industry standards
D. By providing privacy reviews for critical customer applications

Answer: C

Explanation:
Google Cloud commitment to keep the data secure and private:
* 1. Org owns the data and not Google
* 2. Google does not sell data to 3rd parties
* 3. All customer data is encrypted by default
* 4. Google Cloud guards insider against your data
* 5. No backdoor access to any govt. entity
* 6. Google's privacy practices are audited against international standards

NEW QUESTION 60
- (Topic 3)
An organization wants to add a new function to their application. They want to write the code and let the public cloud provider handle the infrastructure.
Which infrastructure solution should they use?

A. Virtual machines
B. Bare Metal Solution
C. Serverless computing
D. Container Registry

Answer: C

Explanation:
Serverless computing , as public cloud prouder(eg. google) will mange the infra things

NEW QUESTION 62
- (Topic 3)
An organization relies on online seasonal sales for the majority of their annual revenue. Why should the organization use App Engine for their customer app?

A. Automatically adjusts physical inventory in real time
B. Autoscales during peaks in demand
C. Runs maintenance during seasonal sales
D. Recommends the right products to customers

Answer: B

NEW QUESTION 65
- (Topic 3)
What is an example of structured data that a healthcare facility stores in their system?

A. X-ray images
B. Surgery video recordings
C. Blood pressure history
D. Physician-written notes

Answer: C

Explanation:
Physical measures like height, weight, blood pressure, blood type, and stage
of the disease can be recorded numerically and they are structured.

NEW QUESTION 66
- (Topic 3)
How does switching from on-premises to the cloud help organizations gain value over time?

A. They can focus their efforts on solution development
B. They can relax their on-premises data security protocols
C. They can expand their internal application hosting infrastructure
D. They can increase development of data recovery systems

Answer: A

NEW QUESTION 70
- (Topic 3)
An organization needs a platform to create custom end-to-end artificial intelligence models. Which Google Cloud product or service should the organization use?

A. Dataproc
B. Compute Engine
C. Recommendations AI

D. Vertex AI

Answer: D

Explanation:
Recommendations AI enables you to build an end-to-end personalized recommendation system based on state-of-the-art deep learning ML models, without a need for expertise in ML or recommendation systems. With Vertex AI, both AutoML training and custom training are available options. Whichever option you choose for training, you can save models, deploy models, and request predictions with Vertex AI. https://cloud.google.com/vertex-ai

NEW QUESTION 73
- (Topic 3)
Why do organizations often struggle to scale their on-premises application infrastructure?

A. Scaling compute instances could breach compliance and/or regulation
B. Increasing compute capacity is time-consuming and costly
C. Their serverless compute functions struggle to meet the demand
D. Their multi-cloud architecture is complex and expensive

Answer: B

NEW QUESTION 76
- (Topic 3)
How would a global organization benefit from managing their data with Cloud Spanner?

A. Cloud Spanner is optimized for cold storage
B. Cloud Spanner replicates data across regions in real time
C. Cloud Spanner is optimized to ingest unstructured data
D. Cloud Spanner visualizes and analyzes data in real time

Answer: B

Explanation:
Spanner is Google's scalable, multi-version, globally-distributed, and synchronously-replicated database.

NEW QUESTION 79
- (Topic 3)
An organization wants to use all available data to offer predictive suggestions on their website that improve over time.
Which method should the organization use?

A. Data automation
B. Trends analysis
C. Machine learning
D. Multiple regression

Answer: C

NEW QUESTION 82
- (Topic 3)
An organization notices that some of their cloud expenditures are too high. What should the organization do to control costs?

A. Streamline the hardware procurement process to reduce costs.
B. Share cost views with the departments to establish more accountability.
C. Change the cost model from operational expenditure to capital expenditure.
D. Ensure that all could resources are tagged with a single tag.

Answer: B

NEW QUESTION 86
- (Topic 3)
An organization wants to migrate legacy applications currently hosted in their data center to the cloud. The current architecture dictates that each application needs its own operating system (OS) instead of sharing an OS.
Which infrastructure solution should they choose?

A. Virtual machines
B. Open source
C. Serverless computing
D. Containers

Answer: A

Explanation:
Virtual machines - you can install customized OS Containers - about applications
Virtualization enables you to run multiple operating systems on the hardware of a single physical server, while containerization enables you to deploy multiple applications using the same operating system on a single virtual machine or server. Serverless computing would be no OS required and the open source operating system allows the use of code that is freely distributed and available to anyone and for commercial purposes such as Linux and Free BSD.

NEW QUESTION 91
- (Topic 3)
Several departments in an organization are working together on a project. The organization wants to customize access to resources for each department.
What is the quickest and most efficient way to achieve this?

A. By mapping IAM roles to job functions for each department
B. By assigning IAM primitive roles to each employee
C. By applying least-privilege to roles for each employee
D. By creating a single shared service account for all departments

Answer: A

NEW QUESTION 94
- (Topic 3)
A manager wants to review Google Cloud data access among their employees. Who is responsible for defining data access policies?

A. Cloud Identity
B. Google Cloud Customer Care team
C. Their organization's IT team
D. Their organization's end users

Answer: C

Explanation:
Cloud Identity and Access Management (IAM) helps customers to define fine-grained access policies and precisely control access to Google Cloud-hosted data.

NEW QUESTION 98
- (Topic 3)
What does Cloud Logging help an organization do?

A. Analyze live source code and log code updates.
B. Deploy infrastructure as code.
C. Analyze logs and accelerate application troubleshooting.
D. Manage storage of custom VM images.

Answer: C

NEW QUESTION 103
- (Topic 3)
Why should an organization consider the total cost of ownership (TCO) when moving from on-premises to the cloud?

A. To evaluate error budget
B. To understand service level availability
C. To evaluate return on investment
D. To calculate required compute power

Answer: C

NEW QUESTION 108
- (Topic 3)
An organization wants to create a new application in the cloud to replace an existing on- premises application.
Which application modernization approach should the organization use?

A. Move the application to the cloud, and then change it.
B. Change their application, and then move it to the cloud.
C. Invent in greenfield.
D. Invent in brownfield.

Answer: D

Explanation:
This approach carries over as much custom components as possible from the source system and minimizes initial reengineering efforts.

NEW QUESTION 113
- (Topic 3)
An organization meets their service level objective (SLO) of 99 999% ("five nines") How much downtime do their end users experience per year?

A. 5 minutes
B. 500 minutes
C. 5 hours
D. 5 days

Answer: A

NEW QUESTION 114
- (Topic 3)

An employee receives an email from their internet service provider asking for their bank account number and password.
Which cybersecurity threat is this?

A. Ransomware
B. Distributed Denial of Service
C. Spamming
D. Phishing

Answer: D

Explanation:
The difference between spam and phishing is that, while they both may be inbox-clogging nuisances, only one (phishing) is actively aiming to steal login credentials and other sensitive data. Spam is a tactic for hawking goods and services by sending unsolicited emails to bulk lists

NEW QUESTION 118
- (Topic 2)
Your customer has a reporting tool that is only occasionally used by the leadership team. Usage of it is frequent - once a week, once a month, or once the quarter. They want to run this application in a cost-effective manner. What are the compute options available on Google Cloud which would be suitable? (Choose Two answer)

A. Cloud Run
B. Cloud App Engine Standard
C. Compute Engine
D. Kubernetes Engine

Answer: AB

Explanation:
Since the use of the tool is infrequent/intermittent, you can choose to compute options that are serverless. Both Cloud Run and Cloud App Engine Standard are serverless options that can shut down to zero. Since cost-effectiveness is a requirement, this will not cost anything during the periods it is not used.

NEW QUESTION 121
- (Topic 2)
You are storing sensitive information in a Cloud Storage bucket. For legal reasons, you need to be able to record all requests that read any of the stored data. You want to make sure you comply with these requirements. What should you do?

A. Scan the bucket using the Data Loss Prevention API.
B. Enable Data Access audit logs for the Cloud Storage API.
C. Enable the Identity Aware Proxy API on the project.
D. Allow only a single Service Account access to read the data.

Answer: B

Explanation:
Logged information
Your Google Cloud projects contain only the audit logs for resources that are directly within the Cloud project. Other Google Cloud resources, such as folders, organizations, and billing accounts, contain the audit logs for the entity itself.

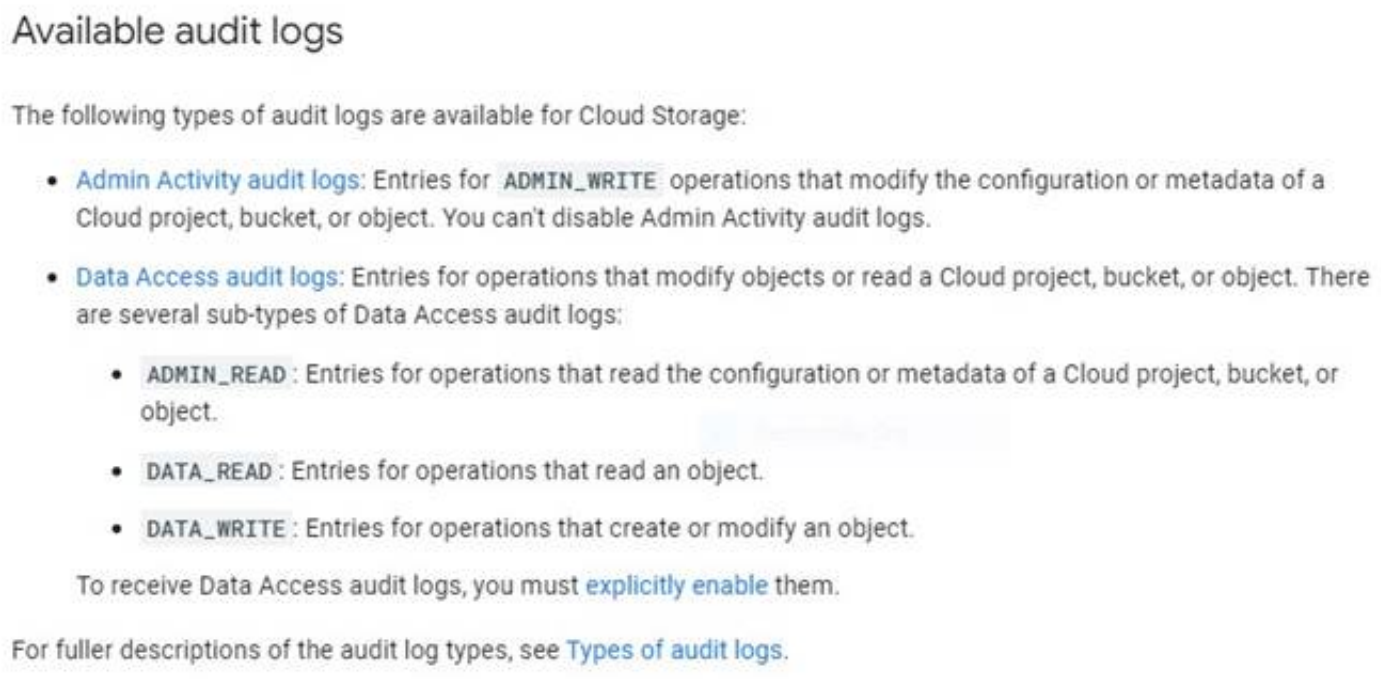
Available audit logs

The following types of audit logs are available for Cloud Storage:

- Admin Activity audit logs: Entries for `ADMIN_WRITE` operations that modify the configuration or metadata of a Cloud project, bucket, or object. You can't disable Admin Activity audit logs.
- Data Access audit logs: Entries for operations that modify objects or read a Cloud project, bucket, or object. There are several sub-types of Data Access audit logs:
  - `ADMIN_READ`: Entries for operations that read the configuration or metadata of a Cloud project, bucket, or object.
  - `DATA_READ`: Entries for operations that read an object.
  - `DATA_WRITE`: Entries for operations that create or modify an object.

To receive Data Access audit logs, you must explicitly enable them.

For fuller descriptions of the audit log types, see Types of audit logs.

Reference link- https://cloud.google.com/storage/docs/audit-logging

NEW QUESTION 125
- (Topic 2)
A large organization is struggling to manage their cloud costs effectively. They want to increase vis-ibility into cloud costs. Which cost management approach should the organization use?

A. Establish a partnership between finance, technology, and business teams.
B. Appoint a single person to monitor cloud spending across the organization.
C. Review any cloud spending that exceeds the organization's error budget.
D. Increase monitoring of on-premises infrastructure and services.

Answer: A

Explanation:
Because cross-team partnerships are part of the visibility cost management strategy.
https://wa.aws.amazon.com/wat.question.COST_1.en.html

NEW QUESTION 130
- (Topic 2)
Google Cloud Platform (GCP) provides three main compliance resource webpages. What are they? (Select Three Answer)

A. Compliance Reports Manager
B. Support Hub
C. Compliance Offerings
D. GDPR Home Page
E. TechCentral

Answer: ACD

Explanation:
Compliance Reports Manager, GDPR Home Page, Compliance Offerings GCP provides three main compliance resource webpages
Compliance Reports Manager
– https://cloud.google.com/security/compliance/compliance-reports-manager

Text, timeline Description automatically generated
Compliance Offerings – https://cloud.google.com/security/compliance/offerings

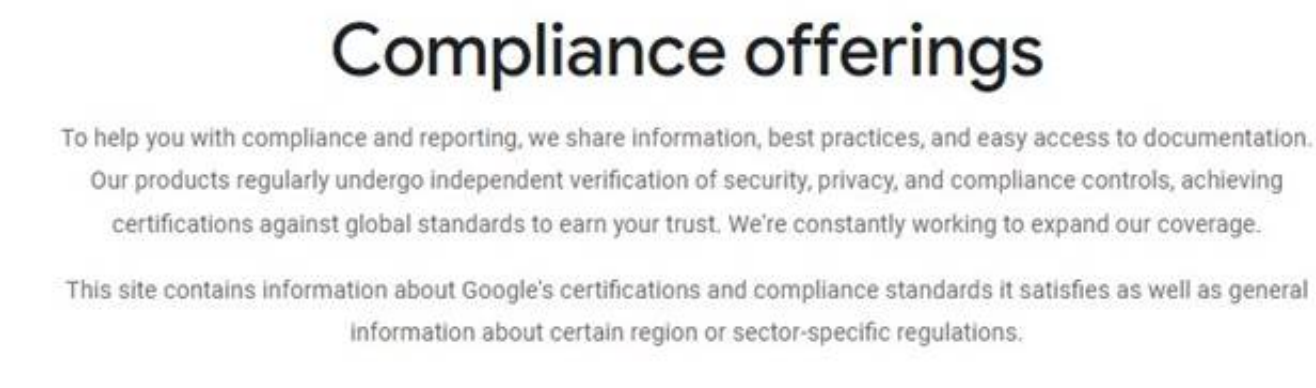

Text Description automatically generated
GDPR Resource Center – https://cloud.google.com/security/gdpr/resource-center At Google Cloud, we champion initiatives that prioritize and improve the security and
privacy of customer personal data, and want you, as a Google Cloud customer, to feel confident using our services in light of GDPR requirements. If you partner with Google Cloud, we will support your GDPR compliance efforts

NEW QUESTION 135
- (Topic 2)
Google offers Firebase, In terms of Firebase Console, any particular message that has to be deliv-ered to a customer at a certain degree of change in behavior can be managed through_____.

A. A/B testing
B. Notification Composer
C. Firebase Remote config.
D. None of the above

Answer: B

Explanation:
You can send notification messages using the Notifications composer in the Firebase console. Though this does not provide the same flexibility or scalability as sending messages with the Admin SDK or the HTTP and XMPP protocols, it can be very useful for testing or for highly targeted marketing and user engagement. The Firebase console provides analytics-based A/B testing to help refine and improve marketing messages.
After you have developed logic in your app to receive messages, you can allow non- technical users to send messages per the instructions on the Notifications

page in the Firebase Help Center.

NEW QUESTION 138
- (Topic 2)
Which of the followings are core components of Anthos?

A. Infrastructure, container, and cluster management
B. Secure software supply chain
C. Multicluster & Configuration management
D. All of the above are correct.

Answer: D

Explanation:

Core Anthos components	Google Cloud	On-premises	Multi-cloud	Attached clusters
Infrastructure, container, and cluster management	GKE Multi Cluster Ingress	Anthos clusters on VMware	Anthos clusters on AWS, Anthos clusters on Azure	
Multicluster management	Fleets, fleet-enabled components, and Connect	Fleets, fleet-enabled components, and Connect	Fleets, fleet-enabled components, and Connect	Fleets, fleet-enabled components, and Connect
Configuration management	Anthos Config Management	Anthos Config Management	Anthos Config Management	Anthos Config Management
Migration	Migrate for Anthos and GKE	Migrate for Anthos and GKE	Migrate for Anthos and GKE	
Service management	Anthos Service Mesh Anthos Service Mesh dashboards MeshCA certificate authority	Anthos Service Mesh Grafana and Kiali dashboards Istiod certificate authority	Anthos Service Mesh (AWS only)	Anthos Service Mesh
Serverless	Cloud Run for Anthos	Cloud Run for Anthos		
Secure software supply chain	Binary Authorization	Binary Authorization (preview)		
Logging and monitoring	Cloud Logging and Cloud Monitoring for system components	Cloud Logging and Cloud Monitoring for system components		
Marketplace	Kubernetes Applications in Cloud Marketplace	Kubernetes Applications in Cloud Marketplace		

NEW QUESTION 140
- (Topic 2)
What are the different storage & database services in GCP? Which is Google cloud storage and da-tabase below the option

A. Persistent Disk
B. Cloud SQL.
C. Cloud Bigtable
D. Cloud Spanner
E. All of the Above

Answer: E

Explanation:

Google Cloud offers 9 storage and database options namely:

- Cloud Storage.
- Cloud SQL.
- Cloud Spanner.
- Cloud Datastore.
- Cloud Bigtable.
- Persistent Disk.
- Cloud Firestore (Firestore & Filestore are both two different types)
- Google Cloud Filestore.

Graphical user interface, text, application Description automatically generated

NEW QUESTION 142
- (Topic 2)
You have contracted a partner to conduct some medical trials. This is a limited, 2-month contract. At the end of each day, you are expecting about 10 Gbs of data. The data is highly sensitive. What networking option would you employ?

A. As the name indicates, set up Partner Interconnect with your partner company.
B. Setup Dedicated Interconnect with your partner.
C. Setup Cloud VPN and create an IPsec VPN tunnel with your partner.
D. Create a public IP for a VM and share that with your partners so that they can access it over the internet and share the data.

Answer: C

Explanation:
"Cloud VPN securely extends your peer network to Google's network through an IPsec VPN tunnel. Traffic is encrypted and travels between the two networks over the public internet. Cloud VPN is useful for low-volume data connections. For additional connection options, see the Hybrid Connectivity product page."

NEW QUESTION 143
- (Topic 2)
A customer has contacted you about migrating to Google Cloud. The customer would like to mi-grate their data from on premises as soon as possible. They don't have the budget to rewrite code, and they want the most direct route. What migration option should suggest to the customer?

A. None, since the customer is not cloud native ready.
B. Rip and Replace
C. Lift and Shift
D. Improve and Move

Answer: C

Explanation:
With Lift and Shift migrations, the customer could move workloads from a source environment to a target environment with few or no modifications or refactoring

Lift and shift

In a lift and shift migration, you move workloads from a source environment to a target environment with minor or no modifications or refactoring. The modifications you apply to the workloads to migrate are only the minimum changes you need to make in order for the workloads to operate in the target environment.

A lift and shift migration is ideal when a workload can operate as-is in the target environment, or when there is little or no business need for change. This migration is the type that requires the least amount of time because the amount of refactoring is kept to a minimum.

There might be technical issues that force a lift and shift migration. If you cannot refactor a workload to migrate and cannot decommission the workload, you must use a lift and shift migration. For example, it can be difficult or impossible to modify the source code of the workload, or the build process isn't straightforward so producing new artifacts after refactoring the source code might not be possible.

Lift and shift migrations are the easiest to perform because your team can continue to use the same set of tools and skills that they were using before. These migrations also support off-the-shelf software. Because you migrate existing workloads with minimal refactoring, lift and shift migrations tend to be the quickest, compared to improve and move or remove and replace migrations.

On the other hand, the results of a lift and shift migration are non-cloud-native workloads running in the target environment. These workloads don't take full advantage of cloud platform features, such as horizontal scalability, fine-grained pricing, and highly managed services.

https://cloud.google.com/architecture/migration-to-gcp-getting-started

NEW QUESTION 144
- (Topic 2)
You are working for a hospital that stores its medical images in an on-premises data room and it is provided that the hospitals want to use Cloud Storage for archival storage of these images. You are required to design and implement a solution where the hospital wants an automated process to up-load any new medical images to Cloud Storage. On the basis of this statements which of the follow-ing statement is correct.

A. Create a Pub/Sub topic, and enable a Cloud Storage trigger for the Pub/Sub topi
B. Create an application that sends all medical images to the Pub/Sub topic.
C. Create a script that uses the gsutil command line interface to synchronize the on-premises storage with Cloud Storag
D. Schedule the script as a cron job.
E. In the Cloud Console, go to Cloud Storag
F. Upload the relevant images to the ap-propriate bucket.
G. Deploy a Dataflow job from the batch template, "Datastore to Cloud Storage" Schedule the batch job on the desired interval.

Answer: B

Explanation:
Using sync for new images implies that you will continue to use your onprem
and keep synchronizing it forever, Sync just once for the old images, new images go directly to google cloud via pub/sub, and eventually get rid of the onprem.

NEW QUESTION 145
- (Topic 2)
Your client has an on-premises data center. Due to technical limitations, they are unable to scale globally. They have decided to adopt the public cloud. However,

they don't want to locked into any one vendor and, therefore, would like to work with multiple cloud providers. They have used open source container technologies and would like to continue using them.

A. Cloud Run which supports containers and can scale in a serverless fashion
B. Kubernetes that runs containers as their core workloads
C. AppEngine Flexible Environment which supports containers
D. Anthos that runs containers as their core workloads

Answer: D

Explanation:
Anthos unifies the management of infrastructure and applications across on- premises, edge, and in multiple public clouds with a Google Cloud-backed control plane for consistent operation at scale.

Anthos enables you to manage GKE clusters and workloads running on virtual machines across environments. You get consistent managed Kubernetes experience with simple installs as well as upgrades validated by Google. Anthos can run on your existing virtualized infrastructure and bare metal servers without a hypervisor layer. Anthos simplifies your application stack, reduces the costs associated with licensing a hypervisor, and decreases time spent learning new skills.

NEW QUESTION 146
- (Topic 2)
Which of these are defined by the following statement: a contract you have with your end custom-ers, which, if you don't meet, you might even have to pay fines?

A. SLA - Service Level Agreement
B. SLC - Service Level Contract
C. SLO - Service Level Objective
D. SLI - Service Level Indicator

Answer: A

Explanation:

Service-Level Agreement (SLA)

At Google, we distinguish between an SLO and a Service-Level Agreement (SLA). An SLA normally involves a promise to someone using your service that its availability SLO should meet a certain level over a certain period, and if it fails to do so then some kind of penalty will be paid. This might be a partial refund of the service subscription fee paid by customers for that period, or additional subscription time added for free. The concept is that going out of SLO is going to hurt the service team, so they will push hard to stay within SLO. If you're charging your customers money, you will probably need an SLA.

Because of this, and because of the principle that availability shouldn't be much better than the SLO, the availability SLO in the SLA is normally a looser objective than the internal availability SLO. This might be expressed in availability numbers: for instance, an availability SLO of 99.9% over one month, with an internal availability SLO of 99.95%. Alternatively, the SLA might only specify a subset of the metrics that make up the internal SLO.

https://cloud.google.com/blog/products/devops-sre/sre-fundamentals-slis-slas-and-slos

NEW QUESTION 149
- (Topic 2)
An organization wants to measure everything as part of its new DevOps philosophy. What should the organization measure?

A. The reliability and health of their systems.
B. The satisfaction and happiness of their employees.
C. The risk and reward of their investments.
D. The speed of their cloud adoption process.

Answer: A

Explanation:
Graphical user interface, text, application, email Description automatically generated

https://newrelic.com/devops/measuring-devops#toc-devops-measurments-for-team-health

NEW QUESTION 151
- (Topic 2)
Considering Google Cloud Storage different Options which of the following is / are correct on the basis of their real world use cases?

A. Cloud Storage : Images, Large Media, files , backups.
B. Google Cloud BigTable : AdTech, Financial and IoT Data.
C. Cloud SQL : User Credentials, customer orders.
D. All of the Above.

Answer: D

Explanation:
Cloud Datastore is the best for semi-structured application data that is used in app engines' applications. Bigtable is best for analytical data with heavy read/write events like AdTech, Financial or IoT data. Cloud Storage is best for structured and unstructured, binary or object data like images, large media files and backups. SQL is best for web frameworks and in existing applications like storing user credentials and customer orders. Cloud Spanner is best for large scale database applications that are larger than two terabytes; for example, for financial trading and e-commerce use cases. As I mentioned at the beginning of the module, depending on your application, you might use one or several of these services to get the job done.

NEW QUESTION 155
- (Topic 2)
You are a program manager in a company and handling a project and you need to create a virtual machine on google cloud console that will be very simple to set up, by flipping a bit via command, API, or with developer console that gives you 30 seconds to shut down when you’re preempted, allow you to save your work that also helps in the company budget upto 70-80% of less charges than the regular VMs.

A. Bare Metal Solutions
B. Preemptible Virtual Machines.
C. Google Cloud VM Instances
D. None of the above.

Answer: B

Explanation:
Preemptible VMs have all these features
Simple configuration
Create a preemptible instance simply by flipping a bit via command, API, or developer console.
Easy extensibility
Attach GPUs and local SSDs to preemptible instances for additional performance and savings.
Graceful shutdown
Compute Engine gives you 30 seconds to shut down when you're preempted, letting you save your work in progress for later.
Large scale computing
Spin up as many instances as you need and turn them off when you're done. You only pay for what you use.
Quickly reclaim capacity
Managed instance groups automatically recreate your instances when they're preempted (if capacity is available).
Fixed pricing
Preemptible VMs have fixed pricing up to 80% off regular instances. They show up on your bill separately so you'll see just how much you're saving.

NEW QUESTION 157
- (Topic 2)
Which of the following is true while creating a boot persistent disk from a snapshot.

A. You cannot apply a snapshot to an existing persistent disk, or apply a snapshot to persistent disks that belong to a different project than that snapshot.
B. It is only possible to apply data from a snapshot when you first create a persistent disk.
C. After you create a snapshot of a boot persistent disk, you can apply data from that snapshot to new persistent disks.
D. All of the above.

Answer: D

Explanation:
When you create a virtual machine (VM) instance, you must also create a boot disk for the VM. You can use a public image, a custom image, or a snapshot that was taken from another boot disk. When you create a boot disk, limit the disk size to 2 TB to account for the limitations of MBR partitioning.
Compute Engine automatically creates a boot persistent disk when you create an instance. If you require additional data storage space for your instances, add one or more secondary instance storage options.
You might need to create a standalone boot persistent disk and attach it to an instance later, or resize a boot persistent disk to improve performance and add more space for additional applications or operating system files. That process is described in Add or resize a persistent disk.
As a best practice, do not use regional persistent disks for boot disks. In a failover situation, they do not force-attach to a VM.

After you create a snapshot of a boot persistent disk, you can apply data from that snapshot to new persistent disks. It is only possible to apply data from a snapshot when you first create a persistent disk. You cannot apply a snapshot to an existing persistent disk, or apply a snapshot to persistent disks that belong to a different project than that snapshot.

NEW QUESTION 159
- (Topic 2)
A client is currently running software on their on-premise systems that is bound by a certain
type of license. They are allowed to run the software on virtualized machines. However, they cannot run them on virtualized machines that are shared by two different companies, teams, or projects. What option do they have on Google Cloud?

A. Google Cloud is a public cloud accessed by multiple customers.
B. Allocate a Bare Metal machine.
C. Setup exclusive login to the VM with self-generated security keys.
D. Allocate sole-tenant nodes

Answer: D

Explanation:
Sole-tenancy lets you have exclusive access to a sole-tenant node, which is a physical Compute Engine server that is dedicated to hosting only your project's VMs. Use sole-tenant nodes to keep your VMs physically separated from VMs in other projects, or to group your VMs together on the same host hardware.
https://cloud.google.com/compute/docs/nodes/sole-tenant-nodes

NEW QUESTION 163
- (Topic 2)
In discussions with a prospective customer who wants to move to Google Cloud to make use of the latest, scalable technologies available therein, you learn that there are very strict regulations concern-ing the storage of data. They only have the approval to store it in their current private data cen-ter. What would you advise them?

A. Retain on-premise itself those portions of data and compute which are under regulatio
B. Take advantage of all the other cloud capabilities for remaining work-loads.
C. It is too risky to touch anything in such a scenari
D. It is best to remain entirely on- premise.
E. Regulations are guideline
F. As long as the data remains encrypted, you can move it anywhere.
G. Petition the government for changes to such regulations as all industries are mov-ing to the public clou
H. Then, when the regulations are eased, move to Google Cloud.

Answer: A

Explanation:
Moving to Google Cloud is not an all-or-nothing option. Certain workloads can continue to remain on-premise while the predominant chunk moves to Google Cloud

NEW QUESTION 166
- (Topic 2)
Your company has a requirement to run manual tests on their web products for UX research before it is released to end customers. The people who will do the tests are external to the company. They will either use their own Gmail id or be given temporary email ids using the applications and record-ing their inputs in another app. The UX testing is done in the last week of the month. Each month the UX testers could be different. How should the IT team manage the users?

A. Since the app is anyways going to be public, create permanent credentials for the UX testers that they can conveniently use each time.
B. It would be a security issue to have users come and g
C. Recommend that the test-ers be permanently hired to plug the vulnerability issue.
D. It would be a security issue to have users come and g
E. Recommend that the test-ers be permanently hired to plug the vulnerability issue.
F. Create a Group with the permissions required to do the test and record their in-put
G. When users arrive each week, add them to the group and after the testing period, remove them from the group.

Answer: D

Explanation:
Groups are convenient to use for this requirement. Permissions to the group are automatically inherited by the members of the group. Adding and removing UX testers from the group will grant and remove permissions.

NEW QUESTION 167
- (Topic 2)
Which of the following statements is/are true about Google Cloud BigTable?

A. It is not compatible with Hadoop.
B. It Scales from Giga Byte to Peta Byte with No Downtime.
C. It can not be used in Real-time Ad analytics and tracking thousands of IoT Devices Data.
D. It is an enterprise-level Database that offers relational and non-relational features

Answer: B

Explanation:
Cloud Bigtable
A fully managed, scalable NoSQL database service for large analytical and operational workloads with up to 99.999% availability.
- Consistent sub-10ms latency—handle millions of requests per second
- Ideal for use cases such as personalization, ad tech, fintech, digital media, and IoT

- Seamlessly scale to match your storage needs; no downtime during reconfiguration
- Designed with a storage engine for machine learning applications leading to better predictions
- Easily connect to Google Cloud services such as BigQuery or the Apache ecosystem

NEW QUESTION 169
- (Topic 2)
A large travel services company has been running all their workloads on Google Cloud in the previous year. They looked at their past usage of cloud resources and see that there is a consistent use of 10,000 virtual machines throughout the year. Based on the projections for the following year they have a strong indication that they will use at least this much or more capacity within Google Cloud. What is one way in which they can take advantage of this knowledge?

A. They can use these numbers to negotiate a better contract with another public cloud number.
B. They can cut costs by cutting down on the number of VMs used.
C. They can get into a committed use contract with Google Cloud to get a significant discount on the usage of VMs.
D. They can ask for a sustained use discount.

Answer: C

Explanation:
Compute Engine lets you purchase committed use contracts in return for deeply discounted prices for VM usage. These discounts are referred to as committed use discounts. Committed use discounts are ideal for workloads with predictable resource needs. When you purchase a committed use contract, you purchase Compute Engine resources—such as vCPUs, memory, GPUs, local SSDs, and sole-tenant nodes—at a discounted price in return for committing to paying for those resources for 1 year or 3 years. The discount is up to 57% for most resources like machine types or GPUs. The discount is up to 70% for memory-optimized machine types.

NEW QUESTION 171
- (Topic 2)
“With cloud messaging you can Customize and deliver messages accordingly to the predetermined time in the user's local time zone.” Comment on the above statement.

A. This statement is undefined.
B. The above statement is partially true.
C. The above statement is completely false.
D. The above statement is completely true.

Answer: D

Explanation:
Firebase Cloud Messaging:
Firebase Cloud Messaging (FCM) is a cross-platform messaging solution that lets you reliably send messages at no cost.
Using FCM, you can notify a client app that new email or other data is available to sync. You can send notification messages to drive user re-engagement and retention. For use cases such as instant messaging, a message can transfer a payload of up to 4000 bytes to a client app.
Key capabilities of Firebase Cloud Messaging:
Send notification messages or data messages: Send notification messages that are displayed to your user. Or send data messages and determine completely what happens in your application code.
Versatile message targeting: Distribute messages to your client app in any of 3 ways—to single devices, to groups of devices, or to devices subscribed to topics.
Send messages from client apps: Send acknowledgments, chats, and other messages from devices back to your server over FCM’s reliable and battery-efficient connection channel.

NEW QUESTION 172
- (Topic 2)
You are a cloud architect in a software solution provider company, one of the client that is a Na-tional Bank who wants to build an application that deals with transactions processing, and it needs a relational database with petabyte of scale data. Which of the following Google Cloud Services will you use?

A. Cloud SQL
B. Cloud Bigtable
C. Cloud Spanner
D. Google Cloud BigQuery

Answer: C

Explanation:
- Cloud Spanner is the online transaction processing solution that is relational and offers petabyte scalability. Cloud SQL is not designed for petabyte-scale data.

NEW QUESTION 176
- (Topic 2)
Cloud SQL is a fully-managed relational database service for MySQL, PostgreSQL and SQL serv-ers, keeping Cloud SQL Google Cloud Service in mind, which of the following statements is/are correct?

A. Data inside cloud SQL is automatically Encrypted.
B. Cloud SQL automatically ensures your databases are reliable, secure, and scala-ble so that your business continues to run without disruption.
C. With DMS (Database Migration Service) it becomes very easy to Migration of Pro- duction Database.
D. All of the above

Answer: D

Explanation:
Cloud SQL
Fully managed relational database service for MySQL, PostgreSQL, and SQL Server. Run the exact same relational databases you know with their rich extension collections, configuration flags and developer ecosystem, but without the hassle of self management.

- Reduce maintenance cost with fully managed MySQL, PostgreSQL and SQL Server databases.
- Ensure business continuity with reliable and secure services backed by 24/7 SRE team.
- Automate database provisioning, storage capacity management, and other time- consuming tasks.
- Database observability made easy for developers with Cloud SQL Insights.
- Easy integration with existing apps and Google Cloud services like GKE and BigQuery.

Key features:
Fully managed
Cloud SQL automatically ensures your databases are reliable, secure, and scalable so that your business continues to run without disruption. Cloud SQL automates all your backups, replication, encryption patches, and capacity increases—while ensuring greater than 99.95% availability, anywhere in the world.
Integrated
Access Cloud SQL instances from just about any application. Easily connect from App Engine, Compute Engine, Google Kubernetes Engine, and your workstation. Open up analytics possibilities by using BigQuery to directly query your Cloud SQL databases. Reliable
Easily configure replication and backups to protect your data. Go further by enabling automatic failover to make your database highly available. Your data is automatically encrypted, and Cloud SQL is SSAE 16, ISO 27001, and PCI DSS compliant and supports HIPAA compliance.
Easy migrations to Cloud SQL
Database Migration Service (DMS) makes it easy to migrate your production databases to Cloud SQL with minimal downtime. This serverless offering eliminates the manual hassle of provisioning, managing, and monitoring migration-specific resources. DMS leverages the
native replication capabilities of MySQL and PostgreSQL to maximize the fidelity and reliability of your migration. And it’s available at no additional charge for native like-to-like migrations to Cloud SQL.

NEW QUESTION 179
- (Topic 2)
A customer is migrating there on-promises data analytics solution to Google Cloud. The current solution has a lot of data being read form and written to disk. The performance of this approach has occasionally been a bottleneck for a scale of operations that your cus- tomer has. The application is fault tolerant and can with stand machine going down fre- quently. In moving to Google Cloud they are asking your advice on any way to improve performance?

A. Use Big Query Which has very fast data access and analysis
B. Use Cloud Storage which can be central, scalable storage
C. Use local SSDs with the VMs
D. Use Persistent Disk with the VMs

Answer: C

Explanation:
Local SSDs are attached to the VM and have very high throughput. However, when the VM shuts down, The local SSD is also shut down, Since our Workload here is foult tolerant, than is not an issue.

NEW QUESTION 182
- (Topic 2)
What cloud deployment model is generally deployed between organizations such as non- profits, hospitals or even enterprises that share similar requirements or interests?

A. Hybrid
B. Community
C. Private
D. Public

Answer: B

Explanation:
Community Cloud – The cloud infrastructure is planned for selective use by a particular community of consumers from organizations that have mutual interests like security needs, policy, and compliance considerations.
Reference link- https://csrc.nist.gov/glossary/term/community_cloud

NEW QUESTION 184
- (Topic 2)
Your customer's IT team is in the process of modernizing their customer-facing applications. They've witnessed others getting good results from employing microservices, and they're keen to adopt it themselves. The first application that they are modernizing has about 5 different sub-parts, which they have identified will be the services. They also identify that each of them has different scale requirements - some services like user login are less frequently used while others like transac-tions are heavily used. What technical strategy would you recommend for them?

A. Containerize the services and orchestrate them with Google Kubernetes Engine.
B. Retain the original application in Compute Engine and scale it as needed using Managed Instance Groups.
C. Retain the original application as a backup and also for separately scaling the ser-vices, create new application binaries.
D. Retain the original application in Compute Engine and scale it as needed using Unmanaged Instance Groups.

Answer: A

Explanation:
Containers and Kubernetes are ideal for the kind of requirement mentioned here - separate microservices that need to scale independently.
Google Kubernetes Engine (GKE) provides a managed environment for deploying, managing, and scaling your containerized applications using Google infrastructure. The GKE environment consists of multiple machines (specifically, Compute Engine instances) grouped together to form a cluster.
Reference link- https://cloud.google.com/kubernetes-engine/docs/concepts/kubernetes- engine-overview

NEW QUESTION 186
- (Topic 2)
You are looking for a one stop reference page for GCP support. What Page would you se- lect?

A. Compliance Hub

B. Google Cloud Platform Status
C. Support Hub
D. Pricing Page

Answer: C

Explanation:
Google provides a page that brings together everything needed around support. Its called the Support Hub
Reference link- https://cloud.google.com/support-hub

NEW QUESTION 190
- (Topic 2)
When creating machine learning models, a key initial step is to identify the type of model required. One of these is the classification model. Which of these statements define a classification model?

A. A type of machine learning model for distinguishing among two or more discrete value
B. E.
C. "book", "car".
D. A type of machine learning model is a meta-model maker, which classifies algo-rithms based on the quality of their output.
E. A type of machine learning model that outputs continuous (typically, floating-point) value
F. E.
G. the predicted price of the house is $120,000.
H. A type of classic model approach that is less used today and which has been re-placed by the regression model.

Answer: A

Explanation:
A classification model classifies the incoming data into one or more discrete classes.

NEW QUESTION 195
- (Topic 2)
Your client's IT environment has so far been on-premises. They run a mix of applications and data-bases on Linux and Windows. They want to move to Google Cloud in the easiest manner possi-ble. What are their best options?

A. Compute Engine with VMs with either Linux or Windows OS.
B. App Engine Standard
C. Cloud Functions
D. Cloud Run

Answer: A

Explanation:
Compute Engine allows you to allocate VMs with different OSs - Windows and Linux, included.

NEW QUESTION 196
- (Topic 2)
What service is a fully managed real-time messaging service that allows you to send and receive messages between independent applications.

A. Cloud Datastore
B. Cloud Pub/Sub
C. Cloud DNS
D. Cloud BigTable
E. Cloud Spanner

Answer: B

Explanation:
Google Cloud Pub/Sub is a scalable, durable event ingestion and delivery system.
-> Pub/Sub allows services to communicate asynchronously, with latencies on the order of 100 milliseconds.
-> Pub/Sub is used for streaming analytics and data integration pipelines to ingest and distribute data. It is equally effective as messaging-oriented middleware for service integration or as a queue to parallelize tasks.
-> Pub/Sub enables you to create systems of event producers and consumers, called publishers and subscribers. Publishers communicate with subscribers asynchronously by broadcasting events, rather than by synchronous remote procedure calls (RPCs).
Reference link- https://cloud.google.com/pubsub/docs/overview

NEW QUESTION 199
- (Topic 2)
What cloud service model would you want to select if you want to solve a particular busi- ness problem by providing CRM services in the cloud to your enterprises?

A. CaaS
B. SaaS
C. PaaS
D. IaaS

Answer: B

Explanation:
SaaS – Software as a Service (SaaS) provides you a complete product that is run and managed by the service provider. You worry only about using the software

and not about infrastructure.
SaaS provides the lowest level of flexibility and management control over the infrastructure. (Example: Google Gsuite and MS O365)

NEW QUESTION 203
- (Topic 2)
You are working in a company where you need to store Terabytes of Image Data daily and process them e.g. Taking photos of the entire planet 24 hours every day with satellite and sending data to data centres to store and process it. Which of the following would be the best combination for your infrastructure.
You are working in a company where you need to store Terabytes of Image Data daily and process them e.g. Taking photos of the entire planet 24 hours every day with satellite and
sending data to data centres to store and process it. Which of the following would be the best combination for your infrastructure.

A. Bare Metal Solutions with Google Cloud Storage.
B. Google Cloud Storage & Google Cloud Compute Engines
C. Google Cloud Storage & Preemptible VMs.
D. None of the Above

Answer: C

Explanation:
The above is a real world example of a company named Planet, where they sent around 80+ satellites to take pictures of earth every day, 24 hours. They run around 40,000 preemptible VMs concurrently.
Preemptible instances function like normal instances but have the following limitations: Compute Engine might stop preemptible instances at any time due to system events. The probability that Compute Engine will stop a preemptible instance for a system event is generally low, but might vary from day to day and from zone to zone depending on current conditions.
Compute Engine always stops preemptible instances after they run for 24 hours. Certain actions reset this 24-hour counter.
Preemptible instances are finite Compute Engine resources, so they might not always be available.
Preemptible instances can't live migrate to a regular VM instance, or be set to automatically restart when there is a maintenance event.
Due to the above limitations, preemptible instances are not covered by any Service Level Agreement (and, for clarity, are excluded from the Compute Engine SLA).
The Google Cloud Free Tier credits for Compute Engine do not apply to preemptible instances.

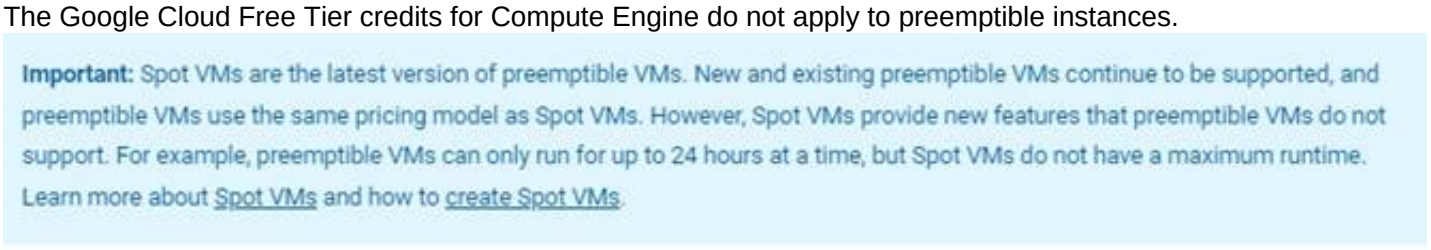

Reference link- https://cloud.google.com/compute/docs/instances/preemptible

NEW QUESTION 207
- (Topic 2)
You are working in a company that provides different services to its customer. Now it also wants to offer some paid API services to its B2B customers for e.g. google provides google maps API, cloud vision API, and language translation API. You need to figure out the best solution for the service.

A. Java Programming Spring Boot Framework for to solve the problem of APIs man-agement.
B. Cloud Functions with Firestore and payment gateways integration development.
C. Apigee API Management
D. Frontend & Backend Development with NodeJs and angular etc.

Answer: C

Explanation:
A top-level idea about Apigee API Management and its offered features can help you solve all questions related to Apigee in Cloud Digital Leader Practice Exam.
Apigee is a platform for developing and managing APIs. By fronting services with a proxy
layer, Apigee provides an abstraction or facade for your backend service APIs and provides security, rate limiting, quotas, analytics, and more.
Apigee services: The APIs that you use to create, manage, and deploy your API proxies. Apigee runtime: A set of containerized runtime services in a Kubernetes cluster that Google maintains. All API traffic passes through and is processed by these services.

NEW QUESTION 211
- (Topic 1)
Your organization needs to build streaming data pipelines. You don’t want to manage the individual servers that do the data processing in the pipelines. Instead, you want a managed service that will automatically scale with the amount of data to be processed.
Which Google Cloud product or feature should your organization choose?

A. Pub/Sub
B. Dataflow
C. Data Catalog
D. Dataprep by Trifacta

Answer: B

Explanation:

Reference: https://cloud.google.com/dataflow/docs/guides/deploying-a-pipeline
Reference link- https://cloud.google.com/dataflow/docs/guides/deploying-a-pipeline

NEW QUESTION 216
- (Topic 1)
The operating systems of some of your organization’s virtual machines may have a security vulnerability.

How can your organization most effectively identify all virtual machines that do not have the latest security update?

A. View the Security Command Center to identify virtual machines running vulnerable disk images
B. View the Compliance Reports Manager to identify and download a recent PCI audit
C. View the Security Command Center to identify virtual machines started more than 2 weeks ago
D. View the Compliance Reports Manager to identify and download a recent SOC 1 audit

Answer: A

Explanation:
Security Health Analytics and Web Security Scanner detectors generate vulnerabilities findings that are available in Security Command Center. Your ability to view and edit findings is determined by the Identity and Access Management (IAM) roles and permissions you are assigned. For more information about IAM roles in Security Command Center.
Reference link:-
https://cloud.google.com/security-command-center/docs/concepts-vulnerabilities-findings

NEW QUESTION 217
- (Topic 1)
The CFO is attending one of the preliminary meetings in the migration strategy meeting. She brings up the concern about costs. They have contracts with their vendors and the payments they will need to make when purchasing any kind of infrastructure. This gives them a clear view of numbers for resource budgeting and planning. Can she get the same kind of clarity on Google Cloud?

A. Ye
B. Do a trial run of typical workload
C. See the billing amount and that becomes the base reference.
D. Yes, the Cloud Native Computing Foundation publishes yearly numbers on the cost of running the clou
E. Use that as a reference.
F. Yes, the Pricing Calculator can be used to estimate the cost of resources.
G. Yes, Google provides a typical cost of application workloads by region and indus-tr
H. Use that as a reference.

Answer: C

Explanation:
The pricing calculator can be used to give clear estimates of resource usage.
-> Running test loads is as closely indicative as using the pricing calculator.
-> There are no cloud cost references published, either by Google or CNCF. Even if some companies have published such info. It might not apply to you.
Reference link:- https://cloud.google.com/products/calculator

NEW QUESTION 221
- (Topic 1)
Your application is onboarding a number of users. The details of the users very widely. What kind of database would be most suitable for this use case?

A. NoSQL database like Firestore
B. OLAP database like BigQuery which support SQL
C. SQL database like MySQL or PostgreSQL
D. OLTP database like Cloud Spanner

Answer: A

Explanation:

* 1. NoSQL databases are best suited for this use case. Firestore is an appropriate one to use here
* 2. Cloud Firestore is a NoSQL document database that lets you easily store, sync, and query data for your mobile and web apps - at global scale.

NEW QUESTION 222
- (Topic 1)
An organization's applications run on an inflexible, on-premises architecture. The organization has decided to modernize their existing applications with the cloud. What may have prompted this business decision?

A. Developers want cloud providers to take full control of their application performance.
B. IT managers want cloud providers to automatically deploy their infrastructure.
C. IT managers want to stop making gradual changes.
D. Developers want to test ideas and experiment with more ease.

Answer: D

Explanation:
Modernizing applications means they can make alterations and innovate more easily.

NEW QUESTION 226
- (Topic 1)
An organization wants to search for and share plug-and-play AI components which can easily build ML services into their project. Which Google Cloud product should the organi- zation use?

A. Document AI
B. AI Hub
C. Cloud Talent Solution
D. Recommendations AI

Answer: B

Explanation:
Because AI Hub is a hosted repository of plug-and-play AI compo-nents. Reference link:- https://cloud.google.com/ai-hub/docs/release-notes

NEW QUESTION 231
- (Topic 1)
Your company needs to segment Google Cloud resources used by each team from the others. The teams' efforts are changing frequently, and you need to reduce operational risk and maintain cost visibility. Which approach does Google recommend?

A. One project per team.
B. One organization per team.
C. One project that contains all of each team's resources.
D. One top-level folder per team.

Answer: A

Explanation:

Reference: https://cloud.google.com/security/infrastructure/design
The Teams need to segmented to have visibility on the resources each team consumes

NEW QUESTION 233
- (Topic 1)
Your company provides car maintenance services. It is conducting an internal hackathon to identify new ideas that could expand their business. The teams have pitched different ideas and have started working on it. They have to present their application to the judges within 48 hours. A presentation alone is not enough; they have to demonstrate a working proof of concept. The team that you are mentoring is going to recommend additional ser-vices to drive in customers based on the brand of car they drive in. They need to be able to identify what brand of car the customer has, based on a photograph automatically taken at entry. They have already discovered an open source database of car images collected by online enthusiasts. How should they implement this solution?

A. Use Deep Learning Containers that are preconfigured and optimized containers for deep learning environments.
B. Use AutoML Image - upload the images and let it create a working model for you.
C. Use TensorFlow to create a model that will identify the car brands; use the available data to train the model.
D. Use Cloud Vision AI that is able to detect logo
E. Write only the code to integrate in-to your workflow.

Answer: B

Explanation:

It would be most straightforward to use AutoML Image. Put the images in Cloud Storage, point to it from AutoML, and start the model building process.
Reference Link- https://cloud.google.com/automl

NEW QUESTION 235
- (Topic 1)
What conditions be true if a VM interface wants to send packets to the external IP addresses of Google APIs and services using Private Google Access?

A. VM interface does not have an external IP address assigned.
B. VM interface is connected to a subnet where Private Google Access is disabled
C. Both A and B
D. None of the Above.

Answer: A

Explanation:
A VM interface can send packets to the external IP addresses of Google APIs and services using Private Google Access if all these conditions are met:
- The VM interface is connected to a subnet where Private Google Access is enabled.
- The VPC network that contains the subnet meets the network requirements for Google APIs and services.
- The VM interface does not have an external IP address assigned.
- The source IP address of packets sent from the VM matches the VM interface's primary internal IP address or an internal IP address from an alias IP range.
A VM with an external IP address assigned to its network interface doesn't need Private Google Access to connect to Google APIs and services. However, the VPC network must meet the requirements for accessing Google APIs and services.

NEW QUESTION 237
- (Topic 1)
Your organization needs a large amount of extra computing power within the next two weeks.
After those two weeks, the need for the additional resources will end. Which is the most cost-effective approach?

A. Use a committed use discount to reserve a very powerful virtual machine
B. Purchase one very powerful physical computer
C. Start a very powerful virtual machine without using a committed use discount
D. Purchase multiple physical computers and scale workload across them

Answer: C

Explanation:
When you purchase a committed use contract, you purchase Compute Engine resources—such as vCPUs, memory, GPUs, local SSDs, and sole-tenant nodes—at a discounted price in return for committing to paying for those resources for 1 year or 3 years

NEW QUESTION 242
- (Topic 1)
A multinational retail company has approached you to help design its systems. They have millions of transactions at their point of sale systems across the world that need to be captured, stored, and analyzed. They are seeing more growth and expect to expand into even more geographies. Which database would be appropriate for them?

A. Cloud Datastore
B. Cloud Storage
C. Cloud Spanner
D. Cloud SQL

Answer: C

Explanation:

Cloud Spanner: "Fully managed relational database with unlimited scale, strong consistency, and up to 99.999% availability."
Reference:- https://cloud.google.com/spanner

NEW QUESTION 247
- (Topic 1)
A prospect wants to be able to store and analyze data. Their analysts already know SQL, but are not familiar with other technologies. Which of these databases can the analysts use without addi-tional training?

A. Cloud SQL, BigQuery, Datastore
B. Spanner, Cloud SQL, BigQuery
C. Cloud SQL, Firestore, Datastore
D. Cloud SQL, Bigtable, BigQuery

Answer: B

Explanation:
Spanner, Cloud SQL, BigQuery
Spanner- Cloud Spanner is a fully managed, mission-critical, relational database service that offers transactional consistency at global scale, automatic, synchronous replication for high availability, and support for two SQL Google Standard SQL and PostgreSQL. Cloud SQL- Cloud SQL is a fully-managed database service that helps you set up,
maintain, manage, and administer your relational databases on Google Cloud Platform. BigQuery- Google BigQuery is a cloud-based Architecture and provides exceptional performance as it can auto-scale up and down based on the data load and performs data analysis efficiently. On the other hand, SQL Server is based on client-server architecture and has fixed performance throughout unless the user scales it manually.

NEW QUESTION 250
- (Topic 1)
Your organization needs to plan its cloud infrastructure expenditures. Which should your organization do?

A. Review cloud resource costs frequently, because costs change often based on use
B. Review cloud resource costs annually as part of planning your organization’s overall budget
C. If your organization uses only cloud resources, infrastructure costs are no longer part of your overall budget
D. Involve fewer people in cloud resource planning than your organization did for on- premises resource planning

Answer: A

Explanation:
Review cloud resource costs frequently, because costs change often based on use because One need to know current usage/ trend for planning; While public cloud eliminates capex, and gets into pay as you go model, the usage pattern determines the cloud cost and hence needs to be measured frequently to enable better forecast

NEW QUESTION 252
- (Topic 1)
Your organization is developing and deploying an application on Google Cloud. Tracking your Google Cloud spending needs to stay as simple as possible. What should you do to ensure that workloads in the development environment are fully isolated from production workloads?

A. Apply a unique tag to development resources
B. Associate the development resources with their own network
C. Associate the development resources with their own billing account
D. Put the development resources in their own project

Answer: D

Explanation:

Reference: https://www.deps.co/blog/google-cloud-platform-good-bad-ugly/
Project resources are components that are necessary for successful project implementation. They include people, equipment, money, time, knowledge – basically, anything that you may require from the project planning to the project delivery phases.

NEW QUESTION 253
- (Topic 1)
A Multiple projects within your organization have long-term VM usage. Based on current demand, they are able to project into the future and estimate how many VM hours they will use every year. Going in for a committed use contract seems sensible. How can you configure committed use easily across all the projects?

A. Enable committed use with discount sharing for the projects
B. Review the usage of resources by each project on a daily basi
C. Enable committed use for the following day based on that number, so that it gives maximum granularity without wastage.
D. Take a report of each project's use in the last yea
E. Enable committed use on a per- project basis based on the previous year's numbers.
F. Share a Google Sheet and request each project team to send in their estimat
G. Enable committed use accordingly on a per-project basis.

Answer: A

Explanation:

Enable committed use discounts are applied to the project from which you purchased it. To share the discount across multiple projects linked to your Cloud Billing account, enable committed use discount sharing from the console. When you enable committed use discount sharing, all of your current active committed use discounts in all the projects under the same Cloud Billing account, including those you previously purchased and new ones you purchase in the future are shared across your Cloud Billing account. Your sustained use discounts are also pooled and shared across all projects within your Cloud Billing account.

NEW QUESTION 254
- (Topic 1)
An organization wants to scale their existing virtual machine architecture as quickly as possible. Why should the organization use VMware Engine?

A. To archive virtual machine instances.
B. To deploy custom APIs seamlessly.
C. To migrate virtual machines to containers.
D. To replatform virtual machines as they are.

Answer: D

Explanation:
VMware Engine helps migrate and run virtual machines in Google Cloud with minimal changes to the VM architecture.

A virtual machine (VM) is a digital version of a physical computer. Virtual machine software can run programs and operating systems, store data, connect to networks, and do other computing functions, and requires maintenance such as updates and system monitoring. Multiple VMs can be hosted on a single physical machine, often a server, and then managed using virtual machine software. This provides flexibility for compute resources (compute, storage, network) to be distributed among VMs as needed, increasing overall efficiency. This architecture provides the basic building blocks for the advanced virtualized resources we use today, including cloud computing.

Learn about virtual machines and VM family types that are available with Compute Engine, the cloud-based computing infrastructure from Google Cloud.

Table
Description automatically generated with medium confidence https://cloud.google.com/learn/what-is-a-virtual-machine

NEW QUESTION 257
- (Topic 1)
A customer has new applications to build that has to handle both batch data and streaming data. Which product should they choose?

A. Dataprep
B. Dataflow
C. Dataproc
D. Data Fusion

Answer: B

Explanation:
Enabling Requester Pays is useful, for example, if you have a lot of data you want to make available to users, but you don't want to be charged for their access to that data.
Reference link- https://cloud.google.com/storage/docs/requester-pays

NEW QUESTION 258
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

Cloud-Digital-Leader Practice Exam Features:

* Cloud-Digital-Leader Questions and Answers Updated Frequently

* Cloud-Digital-Leader Practice Questions Verified by Expert Senior Certified Staff

* Cloud-Digital-Leader Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* Cloud-Digital-Leader Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year