

Exam Questions Cloud-Digital-Leader

Google Cloud Digital Leader exam

<https://www.2passeasy.com/dumps/Cloud-Digital-Leader/>



NEW QUESTION 1

- (Topic 1)

You are leading projects in an IT services company. Your customer's project requires analyzing images. They have many 10s of 1000s of raw images that they have made available to you. Your small technology team needs to build a machine learning model. The images are unlabeled. You don't have the people or the capacity to label the images. What is your approach?

- A. Look for open-source labeled images that closely resemble the given images.
- B. Request data labeling service from Google.
- C. Tell the customer it is their duty to label the images.
- D. Hire temporary workers who can quickly label the images.

Answer: C

Explanation:

Google's Data Labeling Service lets you work with human labelers to generate highly accurate labels for a collection of data that you can use in machine learning models.

References:

-> <https://cloud.google.com/vertex-ai/docs/datasets/data-labeling-job>

-> <https://cloud.google.com/ai-platform/data-labeling/docs>

NEW QUESTION 2

- (Topic 1)

Your organization meant to purchase a 3-year Committed Use Discount, but accidentally purchased a 1-year Committed Use Discount instead. What should your organization do?

- A. Contact your financial institution.
- B. Contact Trust and Safety.
- C. Contact Cloud Billing Support.
- C. Contact Technical Support.

Answer: C

Explanation:

★ **Important:** Once you enable discount sharing through the console, you can only disable it with the assistance of **Cloud Billing support**. If you disable discount sharing with the assistance of Cloud Billing support, all committed use discounts revert to the default setting of applying only to the projects through which they were purchased. The reverted setting becomes effective at the beginning of the following month.

Combining reservations with commitments

A committed use discount provides a 1- or 3-year discounted price agreement, but it does not reserve capacity in a specific zone. A reservation ensures that capacity is held in a specific zone even if the reserved VMs are not running. By combining a reservation with a commitment, you get discounted, reserved resources.

<https://cloud.google.com/compute/docs/instances/signing-up-committed-use-discounts>

NEW QUESTION 3

- (Topic 1)

Your organization runs all its workloads on Compute Engine virtual machine instances. Your organization has a security requirement: the virtual machines are not allowed to access the public internet. The workloads running on those virtual machines need to access BigQuery and Cloud Storage, using their publicly accessible interfaces, without violating the security requirement.

Which Google Cloud product or feature should your organization use?

- A. Identity-Aware Proxy
- B. Cloud NAT (network address translation)
- C. VPC internal load balancers
- D. Private Google Access

Answer: D

Explanation:

VM instances that only have internal IP addresses (no external IP addresses) can use Private Google Access. They can reach the external IP addresses of Google APIs and services. The source IP address of the packet can be the primary internal IP address of the network interface or an address in an alias IP range that is assigned to the interface. If you disable Private Google Access, the VM instances can no longer reach Google APIs and services; they can only send traffic within the VPC network.

Configuring Private Google Access 🔖

[Send feedback](#)

By default, when a Compute Engine VM lacks an external IP address assigned to its network interface, it can only send packets to other internal IP address destinations. You can allow these VMs to connect to the set of external IP addresses used by [Google APIs and services](#) by enabling Private Google Access on the subnet used by the VM's network interface.

Private Google Access also allows access to the external IP addresses used by App Engine, including third-party App Engine-based services.

To view the eligible APIs and services that you can use with Private Google Access, see [supported services](#) in the Private Google Access overview.

See [Private Access Options for Services](#) for background information about Private Google Access and other private connectivity options offered by Google Cloud.

Specifications

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:

Rectangular Strip

- 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 one of the following IP addresses.

If you're sending packets to the [default domains](#):

- The VM interface's primary internal IPv4 address
- The VM interface's internal IPv6 address
- An internal IPv4 address from an alias IP range

<https://cloud.google.com/vpc/docs/configure-private-google-access>

NEW QUESTION 4

- (Topic 1)

Which of the following is/are true about Anthos?

- A. Enterprise-grade container orchestration and management service.
- B. Modernizing your security for hybrid and multi-cloud deployments
- C. Fully managed service mesh with built-in visibility
- D. All of the Above

Answer: D

Explanation:

Anthos :

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.

- Build, deploy, and optimize apps on GKE and VMs anywhere—simply, flexibly, and securely.
- Consistent development and operations experience for hybrid and multi-cloud environments.

Key features:

- * 1. Enterprise-grade container orchestration and management service
- * 2. Automate policy and security at scale
- * 3. Fully managed service mesh with built-in visibility
- * 4. Modernizing your security for hybrid and multi-cloud deployments

NEW QUESTION 5

- (Topic 1)

Which Google Cloud product can report on and maintain compliance on your entire Google Cloud organization to cover multiple projects?

- A. Cloud Logging
- B. Identity and Access Management
- C. Google Cloud Armor
- D. Security Command Center

Answer: D

Explanation:

Security Command Center is a centralized security and risk management platform for your Google Cloud resources. It is a single tool that offers a variety of security features including:

- * 1. Gain centralized visibility and control
- * 2. Discover misconfigurations and vulnerabilities
- * 3. Report on and maintain compliance
- * 4. Detect threats targeting your Google Cloud assets <https://cloud.google.com/security-command-center>

NEW QUESTION 6

- (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 7

- (Topic 1)

Your company has recently acquired three growing startups in three different countries. You want to reduce overhead in infrastructure management and keep your costs low without sacrificing security and quality of service to your customers.

How should you meet these requirements?

- A. Host all your subsidiaries' services on-premises together with your existing services.
- B. Host all your subsidiaries' services together with your existing services on the public cloud.
- C. Build a homogenous infrastructure at each subsidiary, and invest in training their engineers.
- D. Build a homogenous infrastructure at each subsidiary, and invest in hiring more engineers.

Answer: B

Explanation:

Host all your subsidiaries' services together with your existing services on the public cloud.

NEW QUESTION 8

- (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 considerations 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 is the recommended option. References: <https://cloud.google.com/products/operations>

NEW QUESTION 9

- (Topic 1)

Your organization wants to migrate its data management solutions to Google Cloud because it needs to dynamically scale up or down and to run transactional SQL queries against historical data at scale. Which Google Cloud product or service should your organization use?

- A. BigQuery
- B. Cloud Bigtable
- C. Pub/Sub
- D. Cloud Spanner

Answer: D

Explanation:

Reference: <https://cloud.google.com/terms/services>

Cloud Spanner is a fully-managed, mission-critical relational database service. It is designed to provide a scalable online transaction processing (OLTP) database with high availability and strong consistency at global scale

NEW QUESTION 10

- (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 10

- (Topic 1)

An organization wants to dynamically adjust its application to serve different user needs. What are the benefits of storing their data in the cloud for this use case?

- A. Data can be stored in archive for long term access
- B. Automatic data cleaning and validation
- C. Real-time data ingestion and analysis
- D. No data access management required

Answer: C

Explanation:

By storing their application data in the cloud the organization will be able to gather and analyze user behavior data in real-time. This will enable them to dynamically adjust their application for different user needs.

NEW QUESTION 13

- (Topic 1)

Your customer is making a decision on whether to move to Google Cloud. Their key concern is about 10,000 VMs that are part of their IT infrastructure used across more than 110 applications. They are apprehensive of too many changes at this stage. They want to get to Google Cloud in the easiest way possible with minimal disruption. What option would you recommend for them?

- A. Use Migrate for Anthos
- B. Lift and shift the VMs to serverless options like App Engine Flex.
- C. Re-architect on-prem to use Kubernetes and then slowly extend and bridge the on-prem data center to the Google Cloud data center.
- D. Use Migrate for Compute

Answer: D

Explanation:

Migrate for Compute Engine's advanced replication migration technology copies instance data to Google Cloud in the background with no interruptions to the source workload that's running.

Cloud migration creates a lot of questions. Migrate for Compute Engine by Google Cloud has the answers. Whether you're looking to migrate one application from on-premises or one thousand enterprise-grade applications across multiple data centers, Migrate for Compute Engine gives any IT team, large or small, the power to migrate their workloads to Google Cloud.

Watch the video to your right to hear what one of our customers, Rackspace Technology, thinks about Migrate for Compute Engine's speed and ease of use.

<https://cloud.google.com/migrate/compute-engine>

NEW QUESTION 14

- (Topic 1)

Your organization wants to optimize its use of Google Cloud's discounts on virtual machine-based workloads. You plan to use 200 CPUs constantly for the next 3 years, and you forecast that spikes of up to 300 CPUs will occur approximately 30% of the time. What should you choose?

- A. 1-year committed use discount for 200 CPUs
- B. 3-year committed use discount for 300 CPUs
- C. 3-year committed use discount for 200 CPUs
- D. Regular pay-as-you-go pricing

Answer: C

Explanation:

you can get a 57% discount by agreeing to commit to a 3-year contract. Any usage over the commitment will just be billed at the standard rate. Since they only need 300 CPUs 30% of the time, will pick answer C so that we are not paying usage off 300 CPUs all of the time. This gives us a discount of 57% for 200 CPU's, huge cost savings.

NEW QUESTION 15

- (Topic 1)

Your organization wants to be sure that its 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:

Resource hierarchy	Structure and organize your resource hierarchy for fine-grained management and cost allocation using organizations, folders, projects, and labels.
Billing access control	Enforce organizational policies with granular permissions at different levels in the resource hierarchy to control who can spend and who has administrative and cost-viewing permissions.

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 20

- (Topic 1)

A partner of yours used to have their own private data center. Your company was already on Google Cloud and now they have also moved to Google Cloud. You are investigating whether there are ways to collaborate better or shared services. What would be one good option to consider?

- A. Use Private Service Access within Google Cloud.
- B. Use VPC Peering to share resources privately between your two organizations.
- C. Use public IP addresses as before.
- D. It will automatically be routed internally only.
- E. Use VPC Shared Networks to share common resources.

Answer: B

Explanation:

VPC Network Peering allows internal IP address connectivity across two Virtual Private Cloud (VPC) networks regardless of whether they belong to the same project or the same organization.

-> Shared VPC is only within an organization - it allows an organization to connect resources from multiple projects to a common Virtual Private Cloud (VPC) network, so that they can communicate with each other securely and efficiently using internal IPs from that network.

-> Private Google Access is only to access Google APIs and services

References:

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

-> <https://cloud.google.com/vpc/docs/private-google-access>

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

NEW QUESTION 23

- (Topic 1)

A video game organization has invested in cloud technology to generate insights from user behaviors. They want to ensure recommendations of games are aligned to players' interests. What may have prompted this business decision?

- A. Customers expect faster time to market for games.
- B. Employees expect source code changes to be deployed faster.
- C. Customers expect a personalized experience.
- D. Employees expect more predictable data management spending.

Answer: C

Explanation:

Because in the cloud era, users expect more personalization and customization.

NEW QUESTION 28

- (Topic 1)

Which Google Cloud service or feature lets you build machine learning models using Standard SQL and data in a data warehouse?

- A. BigQuery ML
- B. TensorFlow
- C. AutoML Tables
- D. Cloud Bigtable ML

Answer: A

Explanation:

BigQuery ML lets you create and execute machine learning models in BigQuery using standard SQL queries.

Reference: <https://cloud.google.com/bigquery-ml/docs/introduction#:~:text=BigQuery%20ML%20lets%20you%20create,the%20need%20to%20move%20data>

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

<https://cloud.google.com/bigquery-ml/docs/introduction>

NEW QUESTION 29

- (Topic 3)

An organization needs to categorize a large group of photographs using pre-trained machine learning.

Which Google Cloud product or service should the organization use?

- A. Vision API
- B. BigQuery ML
- C. AutoML Vision
- D. Looker

Answer: A

Explanation:

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

NEW QUESTION 34

- (Topic 3)

A large retail organization uses traditional technology for their ecommerce website. During peaks in 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 for 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 39

- (Topic 3)

How is privacy defined in the context of cloud technology?

- A. Restrictions on data access and sharing
- B. Procedures to authenticate user identity
- C. Susceptibility to data breaches and cyber attacks
- D. Compliance with regulatory standards

Answer: A

NEW QUESTION 43

- (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 policies of the country in which the head office is registered

Answer: A

NEW QUESTION 44

- (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 45

- (Topic 3)

An organization needs to migrate specialized workloads to the cloud while maintaining their existing complex licensing and architecture.

What Google Cloud solution should the organization use?

- A. Compute Engine
- B. Bare Metal Solution
- C. Cloud Run
- D. Cloud Functions

Answer: B

Explanation:

“This solution provides a path to modernize your application infrastructure landscape, while maintaining your existing investments and architecture. With Bare Metal Solution, you can bring your specialized workloads to Google Cloud, allowing you access and integration with GCP services with minimal latency.”

NEW QUESTION 48

- (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 52

- (Topic 3)

An organization decides to migrate their on-premises environment to the cloud. They need to determine which resource components still need to be assigned ownership.

Which two functions does a public cloud provider own? (Choose 2) Choose 2 answers

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

Answer: CD

NEW QUESTION 53

- (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 54

- (Topic 3)

An organization finds that the amount of cash in their vending machines doesn't match the value of items sold. They have decided to upgrade their vending machines with cloud-based mobile payment systems.

How could the organization benefit from this upgrade?

- A. They could relax data access permissions.
- B. They could reduce their error budget overspend.
- C. They could improve their perimeter security.
- D. They could view data history to see transactions.

Answer: D

NEW QUESTION 57

- (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 59

- (Topic 3)

An organization wants to search hundreds of scanned documents for key information like dates, names, and other specific words. Why should the organization use application programming interfaces (APIs)?

- A. To replace the scanned documents with an online survey
- B. To ingest data in real time and encrypt unmatched words
- C. To create digital versions of the documents and locate key information
- D. To transform the documents into unstructured data.

Answer: D

Explanation:

The text from the PDF/scanned documents/images gets converted into JSON (unstructured file) which will be further used for search.

NEW QUESTION 61

- (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)

A global organization is developing an application to manage payments and online bank accounts in multiple regions. Each transaction must be handled consistently in their database, and they anticipate almost unlimited growth in the amount of data stored. Which Google Cloud product should the organization choose?

- A. Cloud SQL
- B. Cloud Spanner
- C. Cloud Storage
- D. BigQuery

Answer: B

NEW QUESTION 70

- (Topic 3)

An organization is training a machine learning model to make predictions. What could improve the prediction accuracy of their model?

- A. An increase in storage capacity
- B. Higher network bandwidth
- C. An increase in training data
- D. Faster CPU processors

Answer: C

NEW QUESTION 75

- (Topic 3)

An organization is struggling to meet user demand for change and wants to modernize their legacy applications by moving the applications to the cloud. Why would this help the organization satisfy user expectations?

- A. Toil automation helps make automatic updates
- B. Updates can be pushed out more quickly to repair bugs
- C. Customer data can be used to offer tailored content
- D. DevOps requires that industry trends be measured and tracked

Answer: B

Explanation:

Moving legacy applications to the cloud can help organizations satisfy user expectations by enabling them to push out updates more quickly to repair bugs.

NEW QUESTION 78

- (Topic 3)

An online retail organization wants to optimize their service.

What is an example of unstructured data that they can use to make decisions?

- A. Customer survey comments
- B. Seller location coordinates
- C. Product sales trends
- D. Warehouse inventory records

Answer: A

Explanation:

<https://cloud.google.com/storage/docs/requester-pays>

NEW QUESTION 80

- (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 83

- (Topic 3)

An organization is moving away from an on-premises infrastructure. Instead, they want to create, access, and share information virtually in the cloud.

What should the organization consider?

- A. Built-in security when moving their data to the cloud
- B. Replacing their perimeter security with data encryption keys
- C. Optimizing cost-management with a capital expenditure model
- D. Increased hardware capacity when moving their data to the cloud

Answer: A

NEW QUESTION 88

- (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 92

- (Topic 3)

An organization recently launched a virtual customer support agent, generating vast amounts of text and speech data.

Why should they use a cloud data warehouse to interpret this data?

- A. To natively visualize both types of data using a dashboard in real time
- B. To ingest and analyze structured and unstructured data at scale, in real time
- C. To secure data transmission between cloud and on-premises environments
- D. To transform data from structured to unstructured

Answer: B

Explanation:

Real-time data ingestion and updates. A simple and universal solution for continually ingesting your enterprise data into popular cloud-based data warehouses in real time. <https://www.qlik.com/us/cloud-data-migration/cloud-data-warehouse>

NEW QUESTION 94

- (Topic 3)

A food delivery service needs access to real-time menu information from all partner restaurants. They also need to share customer order information with the

restaurants in real time.
What should the organization use?

- A. Site reliability engineering (SRE)
- B. An application programming interface (API)
- C. A customized machine learning model
- D. A multi-regional database

Answer: B

NEW QUESTION 96

- (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 98

- (Topic 3)

An organization is migrating their business applications from on-premises to the cloud. How could this impact their operations and personnel costs?

- A. Reduced on-premises infrastructure management costs
- B. Increased on-premises hardware maintenance costs
- C. Reduced cloud software licensing costs
- D. Increased cloud hardware management costs

Answer: A

NEW QUESTION 103

- (Topic 3)

An organization wants to write and run small pieces of code in a serverless way that respond to events like huge discounts.
Which Google Cloud compute solution should the organization use?

- A. Google Kubernetes Engine
- B. Cloud Functions
- C. Bare Metal Solution
- D. Compute Engine

Answer: B

NEW QUESTION 106

- (Topic 3)

An organization wants to upskill their IT staff. How can they do this in a transformational way?

- A. Prioritize training current employees instead of hiring new recruits with cloud experience.
- B. Prioritize giving privileged access to third-party partners and contractors to fill IT knowledge gaps.
- C. Create a culture of self-motivated, isolated learning with official training materials.
- D. Create a culture of continuous peer-to-peer learning with official training materials.

Answer: D

NEW QUESTION 111

- (Topic 3)

An organization has created an application that can diagnose different medical conditions when users submit images of their affected body parts.
Which Google Cloud product or service did the organization use?

- A. App Engine
- B. Machine learning
- C. Cloud Logging
- D. Cloud Profiler

Answer: B

NEW QUESTION 112

- (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 116

- (Topic 3)

What is logging within the context of cloud technology?

- A. Writing application and operating system events as text
- B. Monitoring network and resource limitations
- C. Tracking source code across an organization
- D. Recording infrastructure and hardware expenditure

Answer: A

Explanation:

Cloud Logging is a fully managed service that allows you to store, search, analyze, monitor, and alert on logging data and events from Google Cloud and Amazon Web Services

NEW QUESTION 117

- (Topic 3)

An organization needs to run frequent updates for their business app. Why should the organization use Google Kubernetes Engine (GKE)?

- A. Customer expectations can be adjusted without using marketing tools
- B. Seamless changes can be made without causing any application downtime.
- C. GKE handles version control seamlessly and out of the box
- D. GKE is well suited for all monolithic applications

Answer: B

Explanation:

<https://cloud.google.com/architecture/migrating-a-monolithic-app-to-microservices-gke>

NEW QUESTION 122

- (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 124

- (Topic 3)

What is an organization exclusively responsible for when they access an application through a software as a service (SaaS) model?

- A. Maintaining overall system operability
- B. Maintaining customer-facing content
- C. Monitoring data center servers
- D. Monitoring computer networks

Answer: B

NEW QUESTION 125

- (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 126

- (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 128

- (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 132

- (Topic 3)

What is an example of unstructured data that organizations can capture from social media?

- A. Post comments
- B. Tagging
- C. Profile picture
- D. Location

Answer: A

Explanation:

<https://treehousetechgroup.com/8-examples-of-unstructured-data/>

NEW QUESTION 134

- (Topic 2)

You are working with a government agency. A web application serves users of the country. It allows citizens to receive certain services in providing their national identity. Citizens have complained that they are seeing delays in web page loading compared to before. On investigating, they are seeing a lot of spurious traffic coming in from a few IPs which they have identified as foreign. What should they do?

- A. Setup Firewall rules to deny access to the malicious IPs.
- B. Setup Cloud Armor and add the malicious IPs to the deny list.
- C. Setup Firewall rules to allow access only to the IPs from within the country.
- D. Setup Cloud NAT and remove all the internal IPs and replace it with a single public IP.

Answer: B

Explanation:

Cloud Armor provides DDoS protection for applications. It can also "Filter your incoming traffic based on IPv4 and IPv6 addresses or CIDRs. Enforce geography-based access controls to allow or deny traffic based on source geo using Google's geoIP mapping."

NEW QUESTION 135

- (Topic 2)

You are discussing scaling requirements with a gaming company. When the game launches, they are expecting incoming data surges of 2 million users or more during weekends and holidays. Their on-premise systems have had issues scaling and they want your advice on solving the issue. What do you recommend?

- A. Either Compute Engine VMs or Kubernetes nodes work, but it is better to keep a buffer of an extra 2 million users.
- B. We can deploy a Pub/Sub to ingest data which will grow to absorb demand and pass it on to other stages.
- C. We will allocate Compute Engine VMs estimating 80% capacity of 2 million users.
- D. We will allocate Kubernetes nodes estimating 80% capacity of 2 million users.

Answer: B

Explanation:

When there are huge surges in demand, it is preferable to use serverless technologies that automatically scale on demand. In this case, the key concern is data ingestion. Pub/Sub is a serverless system that can expand to absorb such demand.

NEW QUESTION 138

- (Topic 2)

One of your customers used to have a private data center. While within their data center itself, they were consuming some Google services via API calls and other public, well-known addresses published by Google. Now they're evacuating their private data center and are moving to Google Cloud. Could they improve some of their existing architecture with respect to security?

- A. Use VPC Peering with the Google Cloud organization so that you can directly use services using only private IPs.
- B. Use private addresses only
- C. No additional configuration is required
- D. All Google services will be accessible within Google Cloud on private addresses.
- E. Use Shared VPCs with the Google Cloud organization so that you can directly use services using only private IPs.
- F. Enable Private Google Access so that they can remove public IP addresses.

Answer: D

Explanation:

"VM instances that only have internal IP addresses (no external IP addresses) can use Private Google Access. They can reach the external IP addresses of Google APIs and services. If you disable Private Google Access, the VM instances can no longer reach Google APIs and services; they can only send traffic within the VPC network."

Private Google Access

[Send feedback](#)

VM instances that only have internal IP addresses (no external IP addresses) can use Private Google Access. They can reach the external IP addresses of Google APIs and services. The source IP address of the packet can be the primary internal IP address of the network interface or an address in an alias IP range that is assigned to the interface. If you disable Private Google Access, the VM instances can no longer reach Google APIs and services; they can only send traffic within the VPC network.

Private Google Access has no effect on instances that have external IP addresses. Instances with external IP addresses can access the internet, according to the [internet access requirements](#). They don't need any special configuration to send requests to the external IP addresses of Google APIs and services.

You enable Private Google Access on a subnet by subnet basis; it's a setting for subnets in a VPC network. To enable a subnet for Private Google Access and to view the requirements, see [Configuring Private Google Access](#).

<https://cloud.google.com/vpc/docs/private-google-access>

NEW QUESTION 142

- (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 145

- (Topic 2)

A startup client of yours does offline data processing for a few of its clients. They are migrating their applications and the associated data to Google Cloud. They have 100TB of data to move. They presently have a very small private data center setup connected to a local internet provider. The maximum bandwidth they are able to get is 100Mbps. How long will it take them to transfer the data over the internet if the transfer goes smoothly?

- A. About 12 days.
- B. About 2 years.
- C. About 100 days.
- D. About 24 hours.

Answer: C

Explanation:

The key reason I included this question is to clarify some terminologies that will be important for your estimates. The data size mentioned is a TB terabyte. Note the "byte". The speed is mentioned in Mbps, which is Megabits per second. Note the "bits". 8 bits make a byte. So, to get the actual number of bits transferred, you need to multiply the TB number by 8.

Total data transferred (in bits) = $100 * 1,000,000,000,000 * 8$ bits

Speed = 100Mbps = $100 * 1,000,000$. i.e. 100 million bits are transferred per second. Hence time taken to transfer all the data = Total Data/Speed = 8,000,000 seconds.

Number of seconds in a day = $24 * 60 * 60 = 86,400$

Total time taken in days = $8,000,000 / 86,400 = 92.59$ days

Reference link- https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets#online_versus_offline_transfer

NEW QUESTION 147

- (Topic 2)

Which of the following are the current options for paid support in GCP? (Select Three Answer)

- A. Premier
- B. Standard
- C. Enhanced
- D. Role
- E. Premium

Answer: BCE

Explanation:

Because GCP provides three options for paid support which are Standard, Enhanced and Premium.

Basic Support is included with your Google Cloud subscription which cover only Case, phone, and chat support for billing issues only

Reference link- <https://cloud.google.com/support>

NEW QUESTION 148

- (Topic 2)

Which of the following statements is / are correct about Machine Learning?

- A. Machine learning examples include chatbots and automated virtual assistants to automate routine customer service tasks and speed up issue resolution.
- B. Machine learning automates the job of building statistical models with Human Intervention.
- C. Robotic process automation (RPA) can not be attached with ML.
- D. None of the Above.

Answer: A

Explanation:

Customer service

Machine learning examples include chatbots and automated virtual assistants to automate routine customer service tasks and speed up issue resolution.

NEW QUESTION 149

- (Topic 2)

Your company has signed up with a cloud provider and you will be using storage and virtual machines with the provider. The provider has provided your organization some expectations for what the service should perform at. What type of agreement provides a guarantee of a certain level of service such as "Uptime"?

- A. Performance Agreement
- B. Interconnection Agreement
- C. Warranty
- D. Service Level Agreement

Answer: D

Explanation:

Service Level Agreement (SLA)

A service level agreement (SLA) is a contract between a service provider (either internal or external) and the end user that defines the level of service expected from the service provider. Some common SLA's are uptime, Response Time, etc.

NEW QUESTION 150

- (Topic 2)

One of your clients is in the retail sector. They have a small team supporting their operations and a small development team taking care of application development. They have heard of the benefits of machine learning, but they do not have the capacity to hire data scientists or the work to retain them. They have a team of analysts who works primarily on BigQuery and knows how to run SQL queries. They want to be able to get into the new age of machine learning and artificial intelligence. What options are available to run on Google Cloud?

- A. Use the popular open-source libraries SciPy and NumPy to create machine learning models.
- B. Use the Unified AI Platform to create a custom TensorFlow model.
- C. Use BigQuery ML to create machine learning models using SQL queries.
- D. Integrate the Cloud Vision API and the Cloud Speech API to create a custom model that will suit the retail sector.

Answer: C

Explanation:

BigQuery ML allows you to create ML models using standard SQL queries. Those familiar with BigQuery and ML will be able to create ML models with just a basic understanding of machine learning.

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

NEW QUESTION 153

- (Topic 2)

What load balancer type is supported with Cloud Armor security policies?

- A. SSL Proxy, HTTP(S) and SSL
- B. HTTP(S) and SSL
- C. Regional SSL
- D. HTTP(S) Only

Answer: D

Explanation:

Google Cloud Armor security policies protect your application by providing Layer 7 filtering and by scrubbing incoming requests for common web attacks or other Layer 7 attributes to potentially block traffic before it reaches your load balanced backend services or backend buckets. Each security policy is made up of a set of rules that filter traffic based on conditions such as an incoming request's IP address, IP range, region code, or request headers.

-> Google Cloud Armor security policies are available only for backend services behind an external HTTP(S) load balancer. The load balancer can be in Premium Tier or Standard Tier.

-> Google Cloud Armor security policies and IP DENY lists and ALLOW lists are available only for HTTP(S) load balancing.

Reference link- <https://cloud.google.com/armor/docs/security-policy-overview>

NEW QUESTION 155

- (Topic 2)

What is a key difference between VMs and containers?

- A. Virtual Machines take less time to launch; containers take longer to launch.
- B. Virtual Machines can only run Linux; containers can run any operating system.
- C. Virtual Machines use a shared operating system and are therefore lighter; containers are heavier on resources.
- D. Each Virtual Machine in a machine has its own operating system; containers will share the same operating system.

Answer: D

Explanation:

VMs have their individual OSs. All containers on a node use the host operating system.

NEW QUESTION 158

- (Topic 2)

certain devices for cracks, rust, etc. Some of these issues are difficult to identify for a human and your company has seen increasing customer complaints - the customer has paid for an inspection and the field agent said there was no problem, but it later turned out there actually was. The team has come up with a proposal to engage AI to identify issues. On evaluating the existing system, it is seen that the mobile phone network connection is not good or consistent. What solution can work for them?

- A. Use AutoML Vision Edge models.
- B. Use the Rust programming language instead of Python to identify issues like rust.
- C. Use Cloud TPUs which will be able to do the analysis faster on the clou
- D. Thus re-sponses also will be fast.
- E. Use TensorFlow to create custom models and deploy it as TensorFlow Lite mod-els.

Answer: A

Explanation:

AutoML Vision Edge model can be deployed to one of several types of edge devices, such as mobile phones, ARM-based devices, and the Coral Edge TPU
<https://cloud.google.com/vision/automl/docs/edge-quickstart>

NEW QUESTION 160

- (Topic 2)

Which of the following storage options should you use when your company is using Cloud Storage to store application backup files for disaster recovery purposes, provided you want to follow Goog-le's recommended practices.

- A. Multi-Regional Storage
- B. Coldline storage
- C. Nearline Storage
- D. Regional Storage

Answer: B

Explanation:

Coldline storage is a very low cost highly durable storage service for data archiving, online backup, and disaster recovery. Coldline storage is the best choice for data

that you plan to access at most once a year due to its slightly lower availability, 90 day minimum storage duration cost for data access, and higher per operation costs. Nearline and Coldline are for backup and archival storage and having the highest availability for both with 99.9 percent.

NEW QUESTION 164

- (Topic 2)

Your customer is moving from AWS to Google Cloud. Data also needs to be moved. There is about 50TB of data. On AWS, the data resides in an S3 bucket. It is going to be moved to Cloud Storage. Data is also being continuously generated on S3 prior to the cutover. It is preferable that this is also periodically transferred. What is the best way to move the data?

- A. Use the gsutil command-line option
- B. Use the Google Cloud console to drag and drop the files easily
- C. Use the Storage Transfer Service

D. Use a Transfer Appliance

Answer: C

Explanation:

Storage Transfer Service provides options that make data transfers and synchronization easier. We can also schedule one-time transfer operations or recurring transfer operations.

Storage Transfer Service is a product that enables you to:

- Move or backup data to a Cloud Storage bucket either from other cloud storage providers or from a local or cloud POSIX file system.
- Move data from one Cloud Storage bucket to another, so that it is available to different groups of users or applications.
- Move data from Cloud Storage to a local or cloud file system
- Move data between file systems.
- Periodically move data as part of a data processing pipeline or analytical workflow.

Storage Transfer Service provides options that make data transfers and synchronization easier. For example, you can:

- Schedule one-time transfer operations or recurring transfer operations.
- Delete existing objects in the destination bucket if they don't have a corresponding object in the source.
- Delete data source objects after transferring them.
- Schedule periodic synchronization from a data source to a data sink with advanced filters based on file creation dates, filenames, and the times of day you prefer to import data.

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

Reference link- <https://cloud.google.com/architecture/transferring-data-from-amazon-s3-to-cloud-storage-using-vpc-service-controls-and-storage-transfer-service>

NEW QUESTION 167

- (Topic 2)

Keeping Flavours of Apigee in mind, which of the following statements is/are correct?

- A. A hybrid version consisting of a runtime plane installed on-premises or in a cloud provider of your choice, and a management plane running in Apigee's cloud.
- B. In this model, API traffic and data are confined within your own enterprise-approved boundaries.
- C. A hosted SaaS version in which Apigee maintains the environment, allowing you to concentrate on building your services and defining the APIs to those services.
- D. There are two types of Flavours in Apigee i.
- E. Apigee & Apigee Hybrid.
- F. All of the above are correct.

Answer: D

Explanation:

Flavors of Apigee

Apigee comes in the following flavors:

Apigee: A hosted SaaS version in which Apigee maintains the environment, allowing you to concentrate on building your services and defining the APIs to those services.

Apigee hybrid: A hybrid version consisting of a runtime plane installed on-premises or in a cloud provider of your choice, and a management plane running in Apigee's cloud. In this model, API traffic and data are confined within your own enterprise-approved boundaries.

NEW QUESTION 169

- (Topic 2)

In terms of Dockers and Kubernetes, which of the following statements are correct?

- A. Kubernetes uses Docker to deploy, manage, and scale containerized applications.
- B. Difference between Docker and Kubernetes relates to the role each play in containerizing and running your applications
- C. Kubernetes can be used with or without Docker.
- D. All of the above.

Answer: D

Explanation:

Kubernetes vs. Docker

Often misunderstood as a choice between one or the other, Kubernetes and Docker are different yet complementary technologies for running containerized applications.

Docker lets you put everything you need to run your application into a box that can be stored and opened when and where it is required. Once you start boxing up your applications, you need a way to manage them; and that's what Kubernetes does. Kubernetes is a Greek word meaning 'captain' in English. Like the captain is responsible for the safe journey of the ship in the seas, Kubernetes is responsible for carrying and delivering those boxes safely to locations where they can be used.

- Kubernetes can be used with or without Docker.

- Docker is not an alternative to Kubernetes, so it's less of a "Kubernetes vs. Docker" question. It's about using Kubernetes with Docker to containerize your applications and run them at scale.

- The difference between Docker and Kubernetes relates to the role each play in containerizing and running your applications.

- Docker is an open industry standard for packaging and distributing applications in containers.

- Kubernetes uses Docker to deploy, manage, and scale containerized applications.

NEW QUESTION 174

- (Topic 2)

The government has mandated that companies in a particular section of healthcare must retain all the data they collect for a period of 10 years in case an audit needs to be done. Your client, who is in that industry, needs to follow regulations. In addition, your client wants to do an analysis of the data quite frequently in the first year. They also don't want to be liable for any data beyond year 10. What would recommend for your customer?

- A. Use Cloud Storage with nearline storage in year one and Coldline storage thereafter
- B. Use Object lifecycle management to move between storage types and delete them after 10 years.
- C. Use Cloud Storage with standard storage in year one and Coldline storage thereafter
- D. Set a Cloud Scheduler trigger for 1 year to change storage types and 10 years to delete the data.
- E. Use Cloud Storage with standard storage in year one and archival storage thereafter
- F. Use Object lifecycle management to move between storage types and delete them after 10 years.
- G. Use Cloud Storage with standard storage in year one and Coldline storage thereafter
- H. Set a Cloud Tasks to trigger for 1 year to change storage types and 10 years to delete the data.

Answer: C

Explanation:

Cloud storage supports Object Lifecycle Management. To support common

use cases like setting a Time to Live (TTL) for objects, retaining noncurrent versions of objects, or "downgrading" storage classes of objects to help manage costs, Cloud Storage offers the Object Lifecycle Management feature.

Standard storage is recommended for frequently accessed data and Archive for data accessed less than once a year.

Nearline, Coldline, and Archive offer ultra-low-cost, highly-durable, highly available archival storage. For data accessed less than once a year, Archive is a cost-effective storage option for the long-term preservation of data. Coldline is also ideal for cold storage—data your business expects to touch less than once a quarter. For warmer storage, choose Nearline: data you expect to access less than once a month, but possibly multiple times throughout the year.

NEW QUESTION 179

- (Topic 2)

While on-premise, an enterprise had multiple teams, each with its own analytics data store. Attempts to converge the storage for centralized, company-wide analysis failed because of speed and scaling issues. What would be the preferred destination architecture on Google Cloud?

- A. Migrate to Bigtable which provides high throughput reads and writes.
- B. Migrate to Cloud Spanner as a globally scalable SQL database.
- C. Migrate to BigQuery as a central data warehouse.
- D. Migrate to Cloud SQL which supports multiple databases like MySQL, PostgreSQL, and SQL Server - all of the customer's SQL databases can be accommodated here.

Answer: C

Explanation:

BigQuery is the data warehousing option on Google Cloud. Since the source data has already been used for analysis, it should easily fit the BigQuery structure too.

NEW QUESTION 183

- (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

DevOps measurements for reliability and system health

DevOps teams can track system reliability, quality, and overall health using a few key metrics. In DevOps organizations, site reliability engineers, operations engineers, software developers, project managers, and engineering leadership will all find value in these measurements.

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

NEW QUESTION 184

- (Topic 2)

How does a least privilege resource access model contribute to cloud security?

- A. Google is responsible for determining access to cloud resources.
- B. Employees may only access on-premises software with special permission.
- C. Only managers and other senior employees have cloud resource access.
- D. Employees only have access to the cloud resources necessary for their job.

Answer: D

Explanation:

This is the definition of a least privilege model.

A supporting principle that helps organizations achieve these goals is the principle of least privilege. The principle of least privilege addresses access control and states that an individual should have only the minimum access privileges necessary to perform a specific job or task and nothing more

NEW QUESTION 187

- (Topic 2)

A Customer has their current SAP systems using Microsoft SQL Server as the Database. They are migrating to Google Cloud and also preparing to later migrate to the latest version of SAP. The entire IT team is being directed to focus on the migration to the new version of SAP. The new version of SAP does not use Microsoft SQL Server as the Database, Any but the most critical IT management tasks are being deprioritized, How should they migrate their current database to Google Cloud?

- A. Spanner
- B. Bare Metal
- C. BigQuery
- D. Cloud SQL

Answer: D

Explanation:

Cloud SQL supports SQL Server, Since the IT team's attention is being focused on other activities, they will have less time for existing admin tasks, It would be best to take a managed/hosted version.

NEW QUESTION 188

- (Topic 2)

Your Customer's Organization has decided to move to the cloud. They currently run VMs on-premise but their goal on Google cloud is to run containers, primarily on Google Kubernetes Engine. They have a lease for their private data center for another year that they have already paid for. What could be strategy they could adopt in migrating?

- A. Jump and Ramp.
- B. Improve and Move.
- C. Rip and Replace.
- D. Lift and Shift.

Answer: B

Explanation:

Since they have already paid for data center for another year. They have the time and resources to work with, They can make the change to their workloads locally/on- premise Improve and Migrate Move to Google Cloud later on.

NEW QUESTION 193

- (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 198

- (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 203

- (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 207

- (Topic 2)

In Google Cloud IAM: if a policy applied at the project level gives you Owner permissions, your access to an individual resource in that project might be restricted to View permission if someone applies a more restrictive policy directly to that resource. What is correct below the options

- A. False
- B. None of the above.
- C. True
- D. Not defined by GCP.

Answer: A

Explanation:

Policies are a union of those applied to resources themselves and those inherited from higher levels in the hierarchy. If a parent policy is less restrictive, it overrides a more restrictive policy applied to the resource. If a parent policy is more restrictive, it does not override a less restrictive policy applied to the resource. Therefore, access granted at a higher level in the hierarchy cannot be taken away by policies applied at a lower level in the hierarchy.

NEW QUESTION 208

- (Topic 2)

A small scale retailer has been collecting its point of sale transaction in a PostgreSQL Da- tabase. They have raised funding for a strategic expansion goal in the next year that will see them grow significantly in Asia, Europe, North America, Which Database option should they choose in Google Cloud?

- A. BigQuary
- B. Spanner
- C. Cloud SQL
- D. Bigtable

Answer: B

Explanation:

Spanner is a global scale Database that Support SQL querying, Similar to PostgreSQL, Which will be regional. So that will be a fairly smooth move, Since they have the time and the funding, they can plan for this migration.

NEW QUESTION 213

- (Topic 2)

A customer of yours has an SLA with their client that a particular service will respond within 4 sec-onds. The end client has reported that it feels slower. Your engineers do a trial at the client site and notice that there seems to be a delay for many of the requests. It's your team's responsibility to iden-tify the issue quickly within the strict timeline for fixes according to the contract, and then fix it. What should you do?

- A. Recommend a move to serverless technologies which will scale automatically on demand.
- B. Add logging statements at multiple points in the application, build it, and deploy i
- C. Now new requests will give us information on latency in the logs.
- D. Check if the browsers used by the client are different from your
- E. If they are, that's most likely the issu

- F. Ensure that everybody uses the latest version of the browser that you are also using.
G. Use Cloud Trace to collect latency data and track how requests propagate and why there is a delay.

Answer: D

Explanation:

Cloud Trace is a built-in tool in the Operations suite to identify issues like latency.

-> Such fixes are unlikely to change core issues like the service itself being architected or written sub-optimally. Though changes like browser, networking, etc. are helpful, it would be the wrong approach to first recommend that the customer upgrade all their hardware and software.

-> Rewriting code and logging information is going to be time consuming. In general though, logging should always be included in code and it can give good insights. But tracing is way more specific and comprehensive for this requirement.

-> In certain cases, we might identify scaling as the issue. But we should first identify the core problem. So, start with tracing. We can also achieve scale in serverful technologies.

Reference link- <https://cloud.google.com/trace>

NEW QUESTION 214

- (Topic 2)

What characteristics should an organization adopt to be a DevOps organization?

- A. Teamwork over individual work
B. Obsession with Automation over preoccupation with manual work
C. Product based teams over component teams.
D. All of the Above

Answer: D

Explanation:

What characteristics should an organization adopt to be a DevOps organization?

Below are my top 5 characteristics of a DevOps organization.

- Product based teams over component teams. ...
- Obsession with Automation over preoccupation with manual work. ...
- Evidence-based over gut feel. ...
- Teamwork over individual work. ...
- Fail fast over delayed learning.

NEW QUESTION 215

- (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 219

- (Topic 2)

What according to you are NOT the key capabilities of In-App Messaging?

- A. Target messages accordingly to the change in the behavior pattern of the target audience.
B. Creating customized and flexible alerts
C. Increasing conversion for user-to-user sharing
D. Sending relevant messages to the target audience

Answer: C

Explanation:

In-App Messaging

Engage active app users with contextual messages.

Firebase In-App Messaging helps you engage users who are actively using your app by sending them targeted and contextual messages that nudge them to complete key in-app actions - like beating a game level, buying an item, or subscribing to content.

NEW QUESTION 223

- (Topic 2)

Cloud SQL is a fully-managed relational database service for MySQL, PostgreSQL and SQL servers, 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 scalable so that your business continues to run without disruption.
- C. With DMS (Database Migration Service) it becomes very easy to Migration of Production 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 228

- (Topic 2)

Cloud Data Loss Prevention (DLP) is a fully managed service designed to help discover, classify, and protect the most sensitive data. DLP provides three key features (Select Three Answers)

- A. Classification
- B. De-identification
- C. De-classification
- D. Inspection
- E. Reinspection

Answer: ABD

Explanation:

Classification. De-classification and Inspection

Classification is the process to inspect the data and know what data we have, how sensitive it is, and the likelihood. Inspection and classification happen here.

De-identification is the process of removing, masking, replacing information from data.

Reference link- <https://cloud.google.com/dlp/docs>

NEW QUESTION 231

- (Topic 2)

A customer is migrating their on-premises data analytics solution to Google Cloud. The current solution has a lot of data being read from and written to disk. The performance of this approach has occasionally been a bottleneck for a scale of operations that your customer has. The application is fault tolerant and can withstand machine going down frequently. 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 fault tolerant, that is not an issue.

NEW QUESTION 232

- (Topic 2)

With respect to the Core Feature of Standby Instances of Cloud SQL which one of the options is correct.?

- A. The standby instance is used in high availability to replace the primary instance when failover occurs
- B. The standby instance appears in the Google Cloud Console but does not get billed
- C. When failover occurs, connections to the primary instance need to be manually transferred to the standby instance.
- D. The standby instance is used in high availability to replace the primary instance when failover occurs

- E. The standby instance appears in the Google Cloud Console but does not get billed
- F. When failover occurs, connections to the primary instance are automatically transferred to the standby instance.
- G. The standby instance is used in high availability to replace the primary instance when failover occurs
- H. The standby instance doesn't appear in the Google Cloud Console
- I. When failover occurs, connections to the primary instance are automatically transferred to the standby instance.
- J. None of the Above.

Answer: C

Explanation:

The standby instance is used in high availability to replace the primary instance when failover occurs. The standby instance doesn't appear in the Google Cloud Console. When failover occurs, connections to the primary instance are automatically transferred to the standby instance.

Cloud SQL Key Terms: Cloud SQL instance

A Cloud SQL instance corresponds to one virtual machine (VM). The VM includes the database instance and accompanying software containers to keep the database instance up and running.

Database instance

A database instance is the set of software and files that operate the databases: MySQL, PostgreSQL or SQL Server.

High availability

Cloud SQL instances using high availability (HA) provide greater reliability than non-HA instances.

HA in Cloud SQL works by having two synchronized instances: a primary instance and a standby instance. Each instance has exactly one VM. Each instance is in a different zone in the same region.

Failover

A failover is when Cloud SQL switches serving from the original primary instance to the standby instance.

Autofailover is a mechanism that automatically triggers failover when a Cloud SQL instance didn't issue a heartbeat in the previous interval.

Standby instances

The standby instance is used in high availability to replace the primary instance when failover occurs. The standby instance doesn't appear in the Google Cloud Console. When failover occurs, connections to the primary instance are automatically transferred to the standby instance.

Clone

When you clone a Cloud SQL instance, you create a new instance that is a copy of the source instance, but is completely independent. After cloning is complete, changes to the

source instance are not reflected in the clone, and changes in the clone are not reflected in the source instance.

Replication

Replication is the ability to create copies of a Cloud SQL instance or an on-premises database, and offload work to the copies. The main reason for using replication is to scale the use of data in a database without degrading performance on the primary instance. Read replica

The read replica is an exact copy of the primary instance. Data and other changes on the primary instance are updated in almost real time on the read replica.

Send your write transactions to the primary instance, and your read requests to the read replica. The read replica processes queries, read requests, and analytics traffic, thus reducing the load on the primary instance.

Source server

Replication copies transactions from a primary instance to one or more read replicas. The primary instance is also called the source server. The source server can be a Cloud SQL primary instance, or a server outside of Google Cloud, such as an on-premises server or a server running in a different cloud. If the source server is outside of Google Cloud, we call it Replication from an external server.

Cloud SQL Auth proxy client

The Cloud SQL Auth proxy client is open source software maintained by Cloud SQL. It connects to a companion process, the Cloud SQL Auth proxy server, running on your Cloud SQL instance. You run the Cloud SQL Auth proxy client on your own servers. The Cloud SQL Auth proxy client can be used to establish a secure SSL/TLS connection to the database instance, and/or to avoid having to open the firewall. Authentication is done through Identity and Access Management (IAM).

NEW QUESTION 236

- (Topic 2)

Customer Managed Encryption Keys (CMEK) can be used for encrypting data inside Cloud BigTable, which of the following statements is/are correct. (Select two answer)

- A. Administrators can not rotate
- B. Not supported for instances that have clustered in more than one region.
- C. CMEK can only be configured at the cluster level.
- D. You can not use the same CMEK key in multiple projects

Answer: BC

Explanation:

Customer-managed encryption keys for Cloud BigTable.

By default, all the data at rest in Cloud Bigtable is encrypted using Google's default encryption. Bigtable handles and manages this encryption for you without any additional action on your part.

If you have specific compliance or regulatory requirements related to the keys that protect your data, you can use customer-managed encryption keys (CMEK) for BigTable. Instead of Google managing the encryption keys that protect your data, your BigTable instance is protected using a key that you control and manage in Cloud Key Management Service (Cloud KMS).

Features

Security: CMEK provides the same level of security as Google's default encryption but provides more administrative control.

Data access control: Administrators can rotate, manage access to, and disable or destroy the key used to protect data at rest in BigTable .

Auditability: All actions on your CMEK keys are logged and viewable in Cloud Logging. Comparable performance: BigTable CMEK-protected instances offer comparable performance to BigTable instances that use Google default encryption.

Flexibility: You can use the same CMEK key in multiple projects or instances or you can use separate keys, depending on your business needs.

NEW QUESTION 238

- (Topic 2)

The Border Security Agency has hired your software services firm to build an application for them that will collect information about visas stamped on passports. You are given stamped images. You have to find out which country issued the visa and the period of validity. Pull out this data and put it into a database. Which of these applications would be suitable for that?

- A. Use Cloud Vision API - write code to identify the text blocks, copy the data, and store it
- B. Use TensorFlow - write code that will identify the type of visa and the bounding text block

- C. Copy the data and then store it.
- D. Use AutoML - upload other images of visas and run the model creation process which will automatically identify the visas
- E. Use Data Labeling service - outsource the work of marking and extracting the information to others.

Answer: A

Explanation:

Cloud Vision API allows you to programmatically identify images, text, etc. in the document. This would be the best option.
<https://cloud.google.com/vision>

NEW QUESTION 243

- (Topic 2)

You are looking for a one stop reference page for GCP support. What Page would you select?

- 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 244

- (Topic 2)

Your client is a financial services company giving loans based on customer profiles. As part of the regulatory compliance, they have to collect a bunch of different documents with know your customer (KYC) information. They want to be able to process the information in these documents quickly and at scale. They want to integrate the chosen solution as quickly as possible. What are your options on Google Cloud?

- A. Integrate the Cloud Vision API to create a custom model to handle the documents.
- B. Create a model using TensorFlow and integrated it into the process workflow.
- C. Integrate the Lending DocAI and Document AI in two there processes workflow of the processing loan requests.
- D. Integrate the Natural Language API to read the request sent in by clients and to process the forms.

Answer: C

Explanation:

Lending DocAI is a pre-packaged AI solution that speeds "up the mortgage workflow processes to easily process loans and automate document data capture, while ensuring the accuracy and breadth of different documents (e.g., tax statements and asset documents)."
<https://cloud.google.com/solutions/lending-doc-ai>

NEW QUESTION 248

- (Topic 2)

You are a database manager working for a new product that will need millions of reading and writing from the database, with zero downtime, key-value i.e. NoSQL features, no manual steps should be required to ensure consistency, repair data, synchronize writes and deletes, Which of the following database you choose?

- A. Cloud SQL
- B. Cloud BigTable
- C. Cloud Spanner
- D. Cloud Firestore

Answer: B

Explanation:

Cloud BigTable

Key features

High throughput at low latency

Bigtable is ideal for storing very large amounts of data in a key-value store and supports high read and write throughput at low latency for fast access to large amounts of data. Throughput scales linearly—you can increase QPS (queries per second) by adding Bigtable nodes. Bigtable is built with proven infrastructure that powers Google products used by billions such as Search and Maps.

Cluster resizing without downtime

Scale seamlessly from thousands to millions of reads/writes per second. Bigtable throughput can be dynamically adjusted by adding or removing cluster nodes without restarting, meaning 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. Flexible, automated replication to optimize any workload

Write data once and automatically replicate where needed with eventual consistency—giving you control for high availability and isolation of reading and write workloads. No manual steps are needed to ensure consistency, repair data, or synchronize writes and deletes. Benefit from a high availability SLA of 99.999% for instances with multi-cluster routing across 3 or more regions (99.9% for single-cluster instances).

NEW QUESTION 252

- (Topic 2)

A bank wants to track the success of their existing ATM network, which has been modernized with APIs to instantly notify customers about their transfers. What is the benefit of using Apigee to achieve this goal?

- A. It has dashboards that chart dimensions and metrics to report on APIs.
- B. It replicates banking APIs to create new business value.
- C. It measures and tracks their total cost of ownership (TCO).
- D. It allows developers to connect the banking APIs with the public cloud.

Answer: A

Explanation:

Apigee includes analytics services which allow enterprises to report on various aspects of an API.

NEW QUESTION 255

- (Topic 2)

What cloud service model would you want to select if you want to solve a particular business 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 256

- (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.

Important: Spot VMs are the latest version of preemptible VMs. New and existing preemptible VMs continue to be supported, and preemptible VMs use the same pricing model as Spot VMs. However, Spot VMs provide new features that preemptible VMs do not support. For example, preemptible VMs can only run for up to 24 hours at a time, but Spot VMs do not have a maximum runtime. [Learn more about Spot VMs and how to create Spot VMs.](#)

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

NEW QUESTION 259

- (Topic 2)

You have a well established development and operations team. Your teams were managing the entire software delivery/deployment cycle on-premise. When migrating to the cloud, you want to continue having this approach. Which is the ideal option for you?

- A. PaaS - Platform as a Service
- B. SaaS - Software as a Service
- C. IDaaS - Identity as a Service
- D. IaaS - Infrastructure as a Service

Answer: D

Explanation:

IaaS - you're given virtualized resources like VMs, Storage, Network. It is your responsibility to manage everything beyond that. This would be similar to what the organization had on-premise.

NEW QUESTION 261

- (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 262

- (Topic 1)

You are migrating workloads to the cloud. The goal of the migration is to serve customers worldwide as quickly as possible. According to local regulations, certain data is required to be stored in a specific geographic area, and it can be served worldwide. You need to design the architecture and deployment for your workloads.

What should you do?

- A. Select a public cloud provider that is only active in the required geographic area
- B. Select a private cloud provider that globally replicates data storage for fast data access
- C. Select a public cloud provider that guarantees data location in the required geographic area
- D. Select a private cloud provider that is only active in the required geographic area

Answer: C

Explanation:

The goal of the migration is to serve customers worldwide as quickly as possible. According to local regulations, certain data is required to be stored in a specific geographic area, and it can be served worldwide. This characteristic is inherent to the public cloud provider.

NEW QUESTION 263

- (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 264

- (Topic 1)

Your organization is on a critical path with recently developed applications. They are going into production in a month. A few million users are expected to use the new application. They want to ensure minimum disruption when the application goes live. Any issues have to be dealt with within minutes and resolved as quickly as possible. Which Support package should they take?

- A. Enhanced Support
- B. Standard Support
- C. Basic Support
- D. Premium Support

Answer: D

Explanation:

Premium Support will have a 15-minute response time with 24/7 response for high & critical-impact issues.

Premium Support overview

[Send feedback](#)

This page explains the features of Premium Support.

Premium Support is a paid support offering designed for enterprises that run mission critical workloads and require fast response times, platform stability, and increased operational efficiencies.

This overview covers the following aspects of Premium Support:

- How you work with a [Technical Account Manager](#) to optimize your Google Cloud operations and Premium Support experience.
- [Features](#) of the offering, including [support case features](#) and [Customer Aware Support](#).
- [Value Add Services](#) that you can purchase to customize your offering.

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

NEW QUESTION 268

- (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. Yes
- 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 cloud
- 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 industry
- 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 272

- (Topic 1)

Your Google Cloud Platform [GCP] admin has to manage a bunch of API keys for external services that are accessed by different applications, which are used by a few teams. What is the best way to manage them?

- A. Share the information in a Github repository and grant access to the repo in IAM as required.
- B. Store the information in Secret Manager and give IAM read permissions as required.
- C. Store the information in Kubernetes Secrets and only grant read permissions to users as required.
- D. Encrypt the information and store it in Cloud Storage for centralized access
- E. Give the decrypt key only to the users who need to access it.

Answer: B

Explanation:

Store the information in Secret Manager is a secure and convenient storage system for API keys, passwords, certificates, and other sensitive data. Secret Manager provides a central place and single source of truth to manage access, and audit secrets across Google Cloud.

<https://cloud.google.com/secret-manager>

NEW QUESTION 275

- (Topic 1)

Your organization needs to restrict access to a Cloud Storage bucket. Only employees who are based in Canada should be allowed to view the contents. What is the most effective and efficient way to satisfy this requirement?

- A. Deploy the Cloud Storage bucket to a Google Cloud region in Canada
- B. Configure Google Cloud Armor to allow access to the bucket only from IP addresses based in Canada
- C. Give each employee who is based in Canada access to the bucket
- D. Create a group consisting of all Canada-based employees, and give the group access to the bucket

Answer: D

Explanation:

Reference: <https://cloud.google.com/storage/docs/access-control>

Because you can use your own private VPN to access the Canada-only bucket from anywhere in the world.

NEW QUESTION 280

- (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 organization 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 components. Reference link:- <https://cloud.google.com/ai-hub/docs/release-notes>

NEW QUESTION 283

- (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 services 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 286

- (Topic 1)

An organization with hybrid cloud architecture wants to build an application once and be able to run it both on-premises and in their public cloud. Which Google Cloud solution should the organization use?

- A. Cloud Functions
- B. App Engine
- C. Compute Engine
- D. Anthos

Answer: D

Explanation:

Anthos allows organizations to build an application once and run it anywhere.

Migrate directly from VMs, Build, deploy, and optimize apps on GKE, Anthos serverless landing zones and VMs anywhere—simply, flexibly, and securely

A hybrid cloud is one in which applications are running in a combination of different environments. Hybrid cloud computing approaches are widespread because almost no one today relies entirely on the public cloud. Many of you have invested millions of dollars and thousands of hours into on-premises infrastructure over the past few decades. The most common hybrid cloud example is combining a public and private cloud environment, like an on-premises data center, and a public cloud computing environment, like Google Cloud. In the "How-to hybrid" section below, we discuss how some of you may operate a combination of on-premises and multiple public cloud environments, effectively being both hybrid and multicloud.

Want to learn more about Google Cloud's hybrid cloud offering? Check out [Anthos](#).

Reference Link- <https://cloud.google.com/anthos>

NEW QUESTION 288

- (Topic 1)

You are a program manager within a Software as a Service (SaaS) company that offers rendering software for animation studios. Your team needs the ability to allow scenes to be scheduled at will and to be interrupted at any time to restart later. Any individual scene rendering takes less than 12 hours to complete, and

there is no service-level agreement (SLA) for the completion time for all scenes. Results will be stored in a global Cloud Storage bucket. The compute resources are not bound to any single geographical location. This software needs to run on Google Cloud in a cost-optimized way. What should you do?

- A. Deploy the application on Compute Engine using preemptible instances
- B. Develop the application so it can run in an unmanaged instance group
- C. Create a reservation for the minimum number of Compute Engine instances you will use
- D. Start more instances with fewer virtual centralized processing units (vCPUs) instead of fewer instances with more vCPUs

Answer: A

Explanation:

What is a preemptible instance?

Preemptible VM instances are available at much lower price—a 60-91% discount—compared to the price of standard VMs. However, Compute Engine might stop (preempt) these instances if it needs to reclaim the compute capacity for allocation to other VMs. Preemptible instances use excess Compute Engine capacity, so their availability varies with usage.

If your apps are fault-tolerant and can withstand possible instance preemptions, then preemptible instances can reduce your Compute Engine costs significantly. For example, batch processing jobs can run on preemptible instances. If some of those instances stop during processing, the job slows but does not completely stop. Preemptible instances complete your batch processing tasks without placing additional workload on your existing instances and without requiring you to pay full price for additional normal instances.

<https://cloud.google.com/compute/docs/instances/preemptible>

NEW QUESTION 293

- (Topic 1)

An organization has had a data leak scare because one employee made a sensitive Cloud Storage bucket available to the public. Given the nature of the company's business, it is understood that there is never any reason to give the public direct access to any file. The security head wants to ensure that such an event never occurs again. How can you ensure this?

- A. Remove Edit access rights of all Cloud Storage buckets so that no user can make any edits.
- B. Set an organizational policy constraint to restrict bucket access set to the public.
- C. Use Cloud Scheduler to run a job at a specified interval to scan bucket
- D. Any public permissions can be programmatically changed.
- E. Write Cloud Functions code connected to Cloud Storage
- F. Any changes will be notified to the function which can be used to reset the public access.

Answer: B

Explanation:

The straightforward way to set it is using Organizational Policy constraint. Any attempts to change the organizational setting will be rejected for any project and resource.

Introduction to the Organization Policy Service

[Send feedback](#)

The Organization Policy Service gives you centralized and programmatic control over your organization's cloud resources. As the [organization policy administrator](#), you will be able to configure constraints across your entire [resource hierarchy](#).

Benefits

- Centralize control to configure restrictions on how your organization's resources can be used.
- Define and establish guardrails for your development teams to stay within compliance boundaries.
- Help project owners and their teams move quickly without worry of breaking compliance.

References link:

-> <https://cloud.google.com/resource-manager/docs/organization-policy/overview>

-> <https://cloud.google.com/resource-manager/docs/organization-policy/org-policy-constraints>

NEW QUESTION 296

- (Topic 1)

Your organization needs to analyze data in order to gather insights into its daily operations. You only want to pay for the data you store and the queries you perform. Which Google Cloud product should your organization choose for its data analytics warehouse?

- A. Cloud SQL
- B. Dataproc
- C. Cloud Spanner

D. BigQuery

Answer: D

Explanation:

BigQuery is an enterprise data warehouse for large amounts of relational structured data Serverless, highly scalable, and cost-effective multicloud data warehouse designed for business agility.

NEW QUESTION 298

- (Topic 1)

You are currently managing workloads running on Windows Server for which your company owns the licenses. Your workloads are only needed during working hours, which allows you to shut down the instances during the weekend. Your Windows Server licenses are up for renewal in a month, and you want to optimize your license cost.

What should you do?

- A. Renew your licenses for an additional period of 3 year
- B. Renew your licenses for an additional period of 3 year
- C. Negotiate a cost reduction with your current hosting provider wherein infrastructure cost is reduced when workloads are not in use
- D. Renew your licenses for an additional period of 2 year
- E. Negotiate a cost reduction by committing to an automatic renewal of the licenses at the end of the 2 year period
- F. Migrate the workloads to Compute Engine with a bring-your-own-license (BYOL) model
- G. Migrate the workloads to Compute Engine with a pay-as-you-go (PAYG) model

Answer: D

Explanation:

The PAYG model is more convenient because you only pay for usage. And the case describes that the workloads are only run on certain days.

NEW QUESTION 299

- (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 additional 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 303

- (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 308

- (Topic 1)

Which Google Cloud product gives you a consistent platform for multi-cloud application deployments and extends other Google Cloud services to your environment?

- A. Google Kubernetes Engine
- B. Virtual Public Cloud
- C. Compute Engine
- D. Anthos

Answer: D

Explanation:

Anthos

Migrate directly from VMs, Build, deploy, and optimize apps on GKE, Anthos serverless landing zones and VMs anywhere—simply, flexibly, and securely

[Try it free](#)[Contact sales](#)

- ✓ Build, deploy, and optimize apps on GKE and VMs anywhere—simply, flexibly, and securely [Rectangular Snip](#)
- ✓ Consistent development and operations experience for hybrid and multicloud environments
- ✓ Achieve up to 4.8x ROI within 3 years according to the [Forrester Total Economic Impact study](#)
- ✓ Accelerate your VM-based app [migration journey](#) to containers

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

NEW QUESTION 309

- (Topic 1)

Your company is running the majority of its workloads in a co-located data center. The workloads are running on virtual machines (VMs) on top of a hypervisor and use either Linux or Windows server editions. As part of your company's transformation strategy, you need to modernize workloads as much as possible by adopting cloud-native technologies. You need to migrate the workloads into Google Cloud. What should you do?

- A. Export the VMs into VMDK format, and import them into Compute Engine
- B. Export the VMs into VMDK format, and import them into Google Cloud VMware Engine
- C. Migrate the workloads using Migrate for Compute Engine
- D. Migrate the workloads using Migrate for Anthos

Answer: D

Explanation:

Anthos: Anthos lets you build, deploy, and manage applications anywhere in a secure, consistent manner. You can modernize existing applications running on virtual machines while deploying cloud-native apps on containers in an increasingly hybrid and multi-cloud world.

NEW QUESTION 313

- (Topic 1)

You are working with the head of the IT team and planning the move of computing systems. The questionnaire indicates that they have a reporting application that runs almost 24 hours every day of the week. When there is extra load, it queues up the processing and executes tasks when there is less demand. Which of these compute options would you recommend for them?

- A. Use a serverless option - App Engine Standard for Flex
- B. Use a server-based option - Compute Engine.
- C. Use a serverless option - Cloud Functions
- D. Serverless option - Cloud Run

Answer: C

Explanation:

- Because Compute Engine VMs are the preferred compute option as they are long-running.

NEW QUESTION 317

- (Topic 1)

Which of the following statements is/are correct about Bare Metal Solutions?

- A. The network, which Google Cloud manages includes a low-latency Cloud Inter-connect connection into the customer Bare Metal Solution environment.
- B. Bare Metal Solution also includes the provisioning and maintenance of the cus-tom, sole-tenancy hardware with local SAN, and smart hands support.
- C. Bare Metal Solution uses a bring-your-own-license (BYOL) model.
- D. All of the Above.

Answer: D

Explanation:

Option A is true

You are responsible for the licensing of all of your software. Bare Metal Solution uses a bring-your-own-license (BYOL) model.

Apart from this you are responsible for the software, applications, and data that you use and store in the Bare Metal Solution environment.

Responsibilities Data, including:

- Security and encryption
- Backups

Software and applications, including:

- Installation
- Configuration
- Upgrades and patching

Operating system and any hypervisor, including:

- Configuration changes
- Upgrades and patching

- Server clusters, including:
 - Installation
 - Configuration

- Maintenance Licensing

Option B & C is also true.

With Bare Metal Solution, Google Cloud provides and manages the core infrastructure, the network, the physical and network security, and hardware monitoring capabilities in an

environment from which you can access all of the Google Cloud services. The core infrastructure includes secure, controlled-environment facilities, and power.

The Bare Metal Solution also includes the provisioning and maintenance of the custom, sole-tenancy hardware with local SAN, and smart hands support.

The network, which is managed by Google Cloud includes a low-latency Cloud Interconnect connection into the customer Bare Metal Solution environment.

The available Google Cloud services include private API access, management tools, support, and billing.

NEW QUESTION 321

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual Cloud-Digital-Leader Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the Cloud-Digital-Leader Product From:

<https://www.2passeasy.com/dumps/Cloud-Digital-Leader/>

Money Back Guarantee

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 First Try
- * Cloud-Digital-Leader Practice Test Questions in Multiple Choice Formats and Updates for 1 Year