



Google

Exam Questions Cloud-Digital-Leader

Google Cloud Digital Leader exam

NEW QUESTION 1

- (Topic 1)

An organization runs their application on a virtual machine, but every time they want to edit specific features, they have to bring the system offline to update the application. What would be a more appropriate solution for their app?

- A. GPUs
- B. Containers
- C. Hypervisors
- D. Solid State Disk

Answer: B

Explanation:

Because containers can compartmentalize applications which enables parts to be edited in isolation.

What are containers?

Containers are packages of software that contain all of the necessary elements to run in any environment. In this way, containers virtualize the operating system and run anywhere, from a private data center to the public cloud or even on a developer's personal laptop. From Gmail to YouTube to Search, everything at Google runs in containers. Containerization allows our development teams to move fast, deploy software efficiently, and operate at an unprecedented scale. We've learned a lot about running containerized workloads and we've [shared this knowledge](#) with the community along the way: from the early days of contributing [cgroups to the Linux kernel](#), to taking designs from our internal tools and open sourcing them as the [Kubernetes](#) project.

Reference link- <https://cloud.google.com/learn/what-are-containers>

NEW QUESTION 2

- (Topic 1)

Your organization is defining the resource hierarchy for its new application in Google Cloud. You need separate development and production environments. The production environment will be deployed in Compute Engine in two regions. Which structure should your organization choose?

- A. Create a single project for all environment
- B. Use labels to segregate resources by environment.
- C. Create a single project for all environment
- D. Use tags to segregate resources by environment.
- E. Create one project for the development environment and one project for the production environment.
- F. Create two projects for the development environment and two projects for the production environment (one for each region).

Answer: C

Explanation:

Many organizations have separate development and production environments so they can build and test new features without disturbing production traffic. In Optimizely, you can create separate projects for each environment to help with governance. With separate development and production projects, your organization can safely build and QA experiments and Personalization campaigns in a development environment before deploying to production. This approach allows multiple stakeholders in your organization to act as gatekeepers for running new experiments in production.

Set up projects

First, you'll start by creating two new projects: one for development and one for production. Each project will need its own snippet:

1. Create a project for your development environment.
2. Implement the snippet in the head tag for that environment.
3. Add the collaborators who you'd like to have access to your development project.
4. Next, create a project for your production environment.
5. Implement the production project snippet in the head tag of the production environment.
6. Add collaborators who you'd like to have access to your production project.

Reference link- <https://support.optimizely.com/hc/en-us/articles/4410284353805-Set-up-projects-for-development-and-production-environments>

NEW QUESTION 3

- (Topic 1)

Your organization wants an economical solution to store data such as files, graphical images, and videos and to access and share them securely. Which Google Cloud product or service should your organization use?

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

D. BigQuery

Answer: A

Explanation:

- Google Storage is GCP's version of AWS Simple Storage Service (S3) and an S3 bucket would be equivalent to a Google Storage bucket across the two clouds

Despite many external solutions for digital files, some people still store their photos, videos, and content files on their desktop or laptop. The only problem with this method is that your computer can quickly become cluttered with thousands of files. It slows your prized piece of hardware (computer) down.

When you want to find a digital file you probably *expect* that file to come flying up on your screen in an instant. Yet -- anyone who keeps a lot of photos on a computer knows it can take minutes, sometimes hours, to find one - even if you keep it on your desktop. It's just not all that convenient to store things this way. Most importantly, just storing these digital files on a desktop leaves them vulnerable to viruses, damage, or theft. Folks who rely on this also generally don't have a back-up plan.

NEW QUESTION 4

- (Topic 1)

Your organization needs to allow a production job to have access to a BigQuery dataset. The production job is running on a Compute Engine instance that is part of an instance group.

What should be included in the IAM Policy on the BigQuery dataset?

- A. The Compute Engine instance group
- B. The project that owns the Compute Engine instance
- C. The Compute Engine service account
- D. The Compute Engine instance

Answer: C

Explanation:

When an identity calls a Google Cloud API, BigQuery requires that the identity has the appropriate permissions to use the resource. You can grant permissions by granting roles to a user, a group, or a service account.

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

NEW QUESTION 5

- (Topic 1)

Your organization needs to ensure that the Google Cloud resources of each of your departments are segregated from one another. Each department has several environments of its own: development, testing, and production. Which strategy should your organization choose?

- A. Create a project per department, and create a folder per environment in each project.
- B. Create a folder per department, and create a project per environment in each folder.
- C. Create a Cloud Identity domain per department, and create a project per environment in each domain.
- D. Create a Cloud Identity domain per environment, and create a project per department in each domain.

Answer: B

Explanation:

Folders are nodes in the [Cloud Platform Resource Hierarchy](#). A folder can contain projects, other folders, or a combination of both. Organizations can use folders to group projects under the organization node in a hierarchy. For example, your organization might contain multiple departments, each with its own set of Google Cloud resources. Folders allow you to group these resources on a per-department basis. Folders are used to group resources that share common IAM policies. While a folder can contain multiple folders or resources, a given folder or resource can have exactly one parent.

```
# Template for new folder & new project

folder_resource = {
  'name': 'new-folder',
  'type': 'gcp-types/cloudresourcemanager-v2:folders',
  'properties': {
    'parent': 'organizations/99999',
    'displayName': 'new-folder'
  }
}

project_resource = {
  'name': 'new-project',
  'type': 'clouresourcemanager.v1.project',
  'metadata': { 'dependsOn': ['new-folder'] },
  'properties': {
    'name': 'new-project',
    'parent': {
      'type': 'folder',
      # HERE it is -- the problem!
      'id': '${ref.new-folder.name}'
    }
  }
}

return { 'resources': [folder_resource, project_resource] }
```

Rectangular Snip

Reference link- <https://cloud.google.com/resource-manager/docs/creating-managing-folders>Reference link- <https://stackoverflow.com/questions/59460623/how-to-create-a-folder-a-project-under-it-with-deployment-manager-google-cloud>**NEW QUESTION 6**

- (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:

- 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 7

- (Topic 1)

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

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

Answer: D

Explanation:

Cloud Identity Platform

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

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

NEW QUESTION 8

- (Topic 1)

How should a multinational organization that is migrating to Google Cloud consider security and privacy regulations to ensure that it is in compliance with global standards?

- A. Comply with data security and privacy regulations in each geographical region
- B. Comply with regional standards for data security and privacy, because they supersede all international regulations
- C. Comply with international standards for data security and privacy, because they supersede all regional regulations
- D. Comply with regional data security regulations, because they're more complex than privacy standards

Answer: A

Explanation:

Comply with data security and privacy regulations in each geographical region For a multi-national corporation, they need to abide not just by international laws, but also regional laws where they do business.

NEW QUESTION 9

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 10

- (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 10

- (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 13

- (Topic 1)

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

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

Answer: A

Explanation:

Explanation

=> Cloud Pub/Sub- Cloud Pub/Sub is the best to be the end-point for ingesting large amounts of data. It will grow as required, can stream data to downstream systems, and can also work with intermittently available backends.

=> Cloud Dataflow- supports streaming data and therefore is an appropriate option for processing the data that is ingested.

=> BigQuery- BigQuery also supports streaming data and its possible to do real time analytics on it.

=> DataStudio- DataStudio and Looker are for visualization. They don't have any in-built analysis.

=> Cloud Functions- Cloud Functions is a useful serverless endpoint. However, Pub/Sub is better in this case because it can also retain messages for a set period if it was not possible to deliver it first time.

=>Cloud Dataproc- Cloud Dataproc is used for Hadoop/Spark workloads and won't be a good fit here.

NEW QUESTION 14

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 16

- (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 18

- (Topic 1)

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

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

Answer: D

Explanation:

Reference: <https://cloud.google.com/bigtable>

Cloud Bigtable is a sparsely populated table that can scale to billions of rows and thousands of columns, enabling you to store terabytes or even petabytes of data. A single value in each row is indexed; this value is known as the row key. Bigtable is ideal for storing very large amounts of single-keyed data with very low latency. It supports high read and write throughput at low latency, and it is an ideal data source for MapReduce operations.

Bigtable is exposed to applications through multiple client libraries, including a supported extension to the Apache HBase library for Java. As a result, it integrates with the existing Apache ecosystem of open-source Big Data software.

Bigtable's powerful back-end servers offer several key advantages over a self-managed HBase installation:

Incredible scalability. Bigtable scales in direct proportion to the number of machines in

your cluster. A self-managed HBase installation has a design bottleneck that limits the performance after a certain threshold is reached. Bigtable does not have this bottleneck, so you can scale your cluster up to handle more reads and writes.

Simple administration. Bigtable handles upgrades and restarts transparently, and it automatically maintains high data durability. To replicate your data, simply add a second cluster to your instance, and replication starts automatically. No more managing replicas or regions; just design your table schemas, and Bigtable will handle the rest for you.

Cluster resizing without downtime. You can increase the size of a Bigtable cluster for a few hours to handle a large load, then reduce the cluster's size again—all without any downtime. After you change a cluster's size, it typically takes just a few minutes under load for Bigtable to balance performance across all of the nodes in your cluster.

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

NEW QUESTION 23

- (Topic 1)

Your organization is developing a plan for migrating to Google Cloud. What is a best practice when initially configuring your Google Cloud environment?

- A. Create a project via Google Cloud Console per department in your company
- B. Define your resource hierarchy with an organization node on top
- C. Create projects based on team members' requests
- D. Make every member of your company the project owner

Answer: B

Explanation:

The Organization resource is the root node of the Google Cloud resource hierarchy and all resources that belong to an organization are grouped under the organization node. This provides central visibility and control over every resource that belongs to an organization.

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

NEW QUESTION 27

- (Topic 1)

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

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

Answer: AD

Explanation:

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

NEW QUESTION 29

- (Topic 1)

Your large and frequently changing organization's user information is stored in an on-premises LDAP database. The database includes user passwords and group and organization membership.

How should your organization provision Google accounts and groups to access Google Cloud resources?

- A. Replicate the LDAP infrastructure on Compute Engine
- B. Use the Firebase Authentication REST API to create users
- C. Use Google Cloud Directory Sync to create users
- D. Use the Identity Platform REST API to create users

Answer: C

Explanation:

You can run a single instance of Google Cloud Directory Sync to synchronize user accounts and groups to Google Cloud.

Reference: <https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-introduction> Text

Description automatically generated <https://support.google.com/a/answer/106368?hl=en>

NEW QUESTION 32

- (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 35

- (Topic 1)

What are the key features of Google Cloud Identity.

- A. Multi-factor authentication (MFA)
- B. Single sign-on (SSO)
- C. Works with your favorite apps and Endpoint management
- D. All of the Above

Answer: D

Explanation:

Cloud Identity:

A unified identity, access, app, and endpoint management (IAM/EMM) platform.

- Give users easy access to apps with single sign-on.

- Multi-factor authentication protects user and company data.

- Endpoint management enforces policies for personal and corporate devices

KEY FEATURES :

Modernize IT and strengthen security Multi-factor authentication (MFA)

Help protect your user accounts and company data with a wide variety of MFA verification methods such as push notifications, Google Authenticator, phishing-resistant Titan Security Keys, and using your Android or iOS device as a security key.

Endpoint management

Improve your company's device security posture on Android, iOS, and Windows devices using a unified console. Set up devices in minutes and keep your company data more secure with endpoint management. Enforce security policies, wipe company data, deploy apps, view reports, and export details.

Single sign-on (SSO)

Enable employees to work from virtually anywhere, on any device, with single sign-on to thousands of pre-integrated apps, both in the cloud and on-premises.

Works with your favorite apps

Cloud Identity integrates with hundreds of cloud applications out of the box—and we're constantly adding more to the list so you can count on us to be your single identity platform today and in the future.

NEW QUESTION 38

- (Topic 1)

Your organization stores highly sensitive data on-premises that cannot be sent over the public internet. The data must be processed both on-premises and in the cloud.

What should your organization do?

- A. Configure Identity-Aware Proxy (IAP) in your Google Cloud VPC network
- B. Create a Cloud VPN tunnel between Google Cloud and your data center
- C. Order a Partner Interconnect connection with your network provider
- D. Enable Private Google Access in your Google Cloud VPC network

Answer: C

Explanation:

After the service provider provisions your connection, you can start passing traffic between your networks by using the service provider's network.

Reference: <https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview>

NEW QUESTION 39

- (Topic 1)

Your multinational organization has servers running mission-critical workloads on its premises around the world. You want to be able to manage these workloads consistently and centrally, and you want to stop managing infrastructure.

What should your organization do?

- A. Migrate the workloads to a public cloud
- B. Migrate the workloads to a central office building
- C. Migrate the workloads to multiple local co-location facilities
- D. Migrate the workloads to multiple local private clouds

Answer: A

Explanation:

Only public cloud offers to centrally manage the infra. for Pvt cloud it may not be possible to get same Pvt Cloud provider across the globe.

NEW QUESTION 41

- (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 43

- (Topic 1)

What would provide near-unlimited availability of computing resources without requiring your organization to procure and provision new equipment?

- A. Public cloud
- B. Containers
- C. Private cloud
- D. Microservices

Answer: A

Explanation:

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

NEW QUESTION 48

- (Topic 1)

Your organization needs to process large amounts of data from an online application that operates continuously. You do not want to be required to provision infrastructure or create server clusters. What should your organization choose?

- A. Compute Engine with BigQuery
- B. Dataproc
- C. Google Kubernetes Engine with Cloud Bigtable
- D. Dataflow

Answer: D

Explanation:

You do not want to be required to provision infrastructure or create server clusters. Because Unified stream and batch data processing that's serverless, fast, and cost-effective.

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

NEW QUESTION 53

- (Topic 1)

Your company has been using a shared facility for data storage and will be migrating to Google Cloud. One of the internal applications uses Linux custom images that need to be migrated.

Which Google Cloud product should you use to maintain the custom images?

- A. App Engine flexible environment
- B. Compute Engine
- C. App Engine standard environment
- D. Google Kubernetes Engine

Answer: B

Explanation:

Reference: <https://cloud.google.com/compute/docs/images/create-delete-deprecate-private-images>

A custom image is a boot disk image that you own and control access to. Use custom images for the following tasks:

Import a virtual disk to Compute Engine from your on-premises environment or from VMs that are running on your local workstation or on another cloud platform.

You can manually import boot disk images to Compute Engine, but one disk at a time.

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

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

NEW QUESTION 58

- (Topic 1)

The government has ordered an audit of your company's data. You have hired an external company to conduct the audit. They need to be able to review the data stored in your Cloud Storage buckets across eight projects. How would you grant them access?

- A. Give the auditors an Owner role on the eight buckets so that they have proper access.
- B. Give them Storage Object Viewer access to the buckets in those eight projects.
- C. They might need access to all projects as the audit progresses; so give them access to all Storage buckets so that you don't have to do it repeatedly later on.
- D. They might need access to all projects as the audit progresses; so give them the Editor role on all Storage buckets so that you don't have to do it repeatedly later on.

Answer: B

Explanation:

Apply the Principle of Least Privilege and only provide read permissions on only the required buckets. No more, no less

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

NEW QUESTION 60

- (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 64

- (Topic 1)

Which Google Cloud product is designed to reduce the risks of handling personally identifiable information (PII)?

- A. Cloud Storage
- B. Google Cloud Armor
- C. Cloud Data Loss Prevention
- D. Secret Manager

Answer: C

Explanation:

Reference: <https://cloud.google.com/blog/products/gcp/take-charge-of-your-sensitive-data-with-the-cloud-dlp-api>

Cloud Data Loss Prevention: Fully managed service designed to help you discover, classify, and protect your most sensitive data.

NEW QUESTION 67

- (Topic 1)

Your team has developed a machine learning model for your customer. The test results indicate very strong predictive capability. The model is then deployed in production. Evaluation of the predictions in production show that they are off by a pronounced margin. What is the issue and how can you solve for it?

- A. The model is under fitted
- B. Train with less data.
- C. The model is over fitted
- D. Add more features to the model to fix it.
- E. The model is fine since the test results are good
- F. Fix the production of incoming data.
- G. The model is overfitted
- H. Train with more data.

Answer: D

Explanation:

If our ML model does well on the training set than on the production set, then we're likely over fitting. Training with more data would be one solution.

NEW QUESTION 70

- (Topic 3)

How would an organization benefit from using Looker?

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

Answer: D

Explanation:

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

NEW QUESTION 75

- (Topic 3)

An organization needs frequent access to only a subset of their data. They want to reduce costs by depositing the rest of their data across Nearline Coldline and Archive repositories
Which Google Cloud product should the organization use?

- A. Filestore
- B. Cloud Spanner
- C. Data Catalog
- D. Cloud Storage

Answer: D

Explanation:

Per Google docs, specifically for GCP Cloud Storage there exists four types of storage with one of them, standard storage, being described as "storage for data that is frequently accessed ("hot" data) and/or stored for only brief periods of time." <https://cloud.google.com/storage>

NEW QUESTION 79

- (Topic 3)

An organization wants to collect metrics and metadata from their cloud applications and put them into dashboards.
Which Google Cloud tool should they use?

- A. Cloud Monitoring
- B. Cloud Trace
- C. Cloud Logging
- D. Cloud Debugger

Answer: A

Explanation:

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

NEW QUESTION 80

- (Topic 3)

What is monitoring within the context of cloud operations?

- A. Observing cloud expenditure in real time to ensure that budgets are not exceeded
- B. Collecting predefined and custom metrics from applications and infrastructure
- C. Tracking user activities to guarantee compliance with privacy regulations
- D. Tracing user location to document regional access and utilization

Answer: B

NEW QUESTION 85

- (Topic 3)

An organization's web developers and operations personnel use different systems. How will increasing communication between the teams reduce issues caused by silos?

- A. By assigning blame for failures and establishing consequences
- B. By combining job role responsibilities to ensure that everyone has shared access
- C. By increasing data encryption to strengthen workflows
- D. By emphasizing shared ownership of business outcomes

Answer: D

NEW QUESTION 90

- (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 91

- (Topic 3)

An organization is making a strategic change to customer support in response to feedback. They plan to extend their helpline availability hours. Why is the organization making this change?

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

Answer: C

NEW QUESTION 96

- (Topic 3)

An organization is training a machine learning model to predict extreme weather events in their country. How should they collect data to maximize prediction accuracy?

- A. Collect all weather data evenly across all cities
- B. Collect all weather data primarily from at-risk cities
- C. Collect extreme weather data evenly across all cities
- D. Collect extreme weather data primarily from at-risk cities

Answer: A

Explanation:

Collect all weather data evenly across all cities. Mainly because it seems that the emphasis for data collection for ML is to make sure there are no holes in your data collection.

NEW QUESTION 98

- (Topic 3)

An organization is altering their gaming product so that it is compatible with cloud technology. What can they expect when moving from traditional technology to cloud technology?

- A. No change to existing responsibilities
- B. A shift toward OpEx
- C. A shift toward using structured data
- D. Increased hardware maintenance

Answer: B

NEW QUESTION 99

- (Topic 3)

An organization needs to categorize text-based customer reviews on their website using a pre-trained machine learning model. Which Google Cloud product or service should the organization use?

- A. Cloud Natural Language API
- B. Dialogflow
- C. Recommendations AI
- D. TensorFlow

Answer: A

Explanation:

<https://cloud.google.com/natural-language>

Use entity analysis to find and label fields within a document—including emails, chat, and social media—and then sentiment analysis to understand customer opinions to find actionable product and UX insights.

NEW QUESTION 103

- (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 106

- (Topic 3)

After rolling out a new update, an organization found a minor bug in its online video game. How should the organization approach this bug while following SRE principles?

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

Answer: A

Explanation:

<https://www.blameless.com/sre/sre-principles>

Accepting failure as normal is one of the SRE principles. SREs believe that accepting failure as normal helps to build an iterative, collaborative culture. One way this is done is by holding a blameless "lessons learned" discussion after an incident occurs.

NEW QUESTION 108

- (Topic 3)

An organization is struggling to keep up with the growth of their application which is running on legacy infrastructure.

What might be holding them back?

- A. The inaccessibility of their data due to perimeter security
- B. The overreliance on platform as a service
- C. The time it takes their serverless compute function to scale
- D. The cost of provisioning hardware for peak usage

Answer: D

Explanation:

Legacy infrastructure is typically based on on-premises hardware that is managed and maintained by the organization. As the application grows and the user base expands, the hardware required to support it must also grow. This can lead to significant costs associated with provisioning and maintaining hardware, particularly if the organization needs to provision for peak usage.

NEW QUESTION 111

- (Topic 3)

An organization delivers a proactive healthcare service. They want to efficiently and automatically collect patient data.

What should the organization encourage the patients to do?

- A. Use at-home health screening devices and then upload their health data daily
- B. Wear Internet of Things (IoT) devices that upload their health data in real time
- C. Self-assess their health data and then document and upload it in real time
- D. Visit a nurse who will use Internet of Things (IoT) devices to collect and upload their health data

Answer: B

NEW QUESTION 113

- (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 114

- (Topic 3)

An organization is planning its cloud expenditure. What should the organization do to control costs?

- A. Consider cloud resource costs as capital expenditure in annual planning.
- B. Use only cloud resources; they have no cloud infrastructure costs.
- C. Review cloud resource costs frequently because costs depend on usage.
- D. Assess cloud resources costs only when SLO is not met by their cloud provider.

Answer: C

NEW QUESTION 116

- (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 121

- (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 124

- (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 127

- (Topic 3)

A retail company stores their product inventory in a legacy system. Often, customers find products on the company's website and want to purchase them in-store. However, when they arrive, they discover that the products are out of stock.

How could the company benefit from using an application programming interface (API)?

- A. To create personalized product recommendations for customers
- B. To optimize their on-premises legacy system stability
- C. By manually linking each inventory system to the website on a case-by-case basis
- D. By programmatically connecting the inventory system to their website

Answer: D

Explanation:

By programmatically connecting the inventory system to their website The issue is the website shows an item is available at the store, but when the customer gets to the store, they find out that item is out of stock.

NEW QUESTION 132

- (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 133

- (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 135

- (Topic 3)

What DevOps practice should an organization use when developing their application to help minimize disruption caused by bugs?

- A. Pause production until all bugs have been eliminated
- B. Prioritize fixing large bugs during production because they are easier to review
- C. Implement small changes incrementally to reduce recovery time when bugs appear
- D. Implement large changes together to make rolling back easier when bugs appear

Answer: C

Explanation:

One of the key principles of DevOps is to release changes frequently and in small batches. This helps to reduce the risk of disruption caused by bugs. If a bug is introduced in a small change, it is easier to identify and fix the bug without affecting a large number of users.

NEW QUESTION 137

- (Topic 3)

An organization notices that some of their cloud expenditures are too high. What should the organization do to control costs?

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

Answer: B

NEW QUESTION 138

- (Topic 3)

An organization needs to store structured, semi-structured, and unstructured data in its raw, native format in the same repository. Which cloud data management solution should the organization use?

- A. Data field
- B. Data lake
- C. Database
- D. Data warehouse

Answer: B

Explanation:

A data lake can store all types of data with no fixed limitation on account size or file and with no specific purpose defined yet. The data comes from disparate sources and can be structured, semi-structured, or even unstructured. Data-lake data can be queried as needed.

<https://cloud.google.com/learn/what-is-a-data-lake>

A data lake is a centralized repository designed to store, process, and secure large amounts of structured, semistructured, and unstructured data. It can store data in its native format and process any variety of it, ignoring size limits.

NEW QUESTION 142

- (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 147

- (Topic 3)

A manager wants to review Google Cloud data access among their employees. Who is responsible for defining data access policies?

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

Answer: C

Explanation:

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

NEW QUESTION 152

- (Topic 3)

Which technology allows organizations to run multiple computer operating systems on a single piece of physical hardware?

- A. Hypervisor
- B. Containers
- C. Serverless computing
- D. Open source

Answer: A

NEW QUESTION 155

- (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 158

- (Topic 3)

An organization wants to develop an application that can be personalized to user preferences throughout the year. Why should they build a cloud-native application instead of modernizing their existing on-premises application?

- A. Developers can rely on the cloud provider for all source code
- B. Developers can launch new features in an agile way
- C. IT managers can migrate existing application architecture without needing updates
- D. IT managers can accelerate capital expenditure planning

Answer: B

NEW QUESTION 161

- (Topic 3)

An organization wants to move from a tactical cloud adoption approach to a transformational approach. How should they change their cloud security?

- A. Provide staff identities using only Google Cloud authentication.
- B. Provide multiple layers of network security using a zero-trust model.
- C. Emphasize strong perimeter security and trust in their private network.
- D. Emphasize three main Identity Access Management roles: owner, editor, and viewer.

Answer: B

Explanation:

<https://www.crowdstrike.com/cybersecurity-101/zero-trust-security/>

Zero Trust is a security framework requiring all users, whether in or outside the organization's network, to be authenticated, authorized, and continuously validated for security configuration and posture before being granted or keeping access to applications and data

NEW QUESTION 162

- (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 167

- (Topic 3)

An organization cannot afford to modernize their infrastructure but they want to process data from their legacy system in a modern platform hosted by a business partner

What solution should the organization choose to make their data accessible?

- A. Compute Engine
- B. Anthos
- C. An application programming interlace
- D. Google Kubernetes Engine

Answer: C

NEW QUESTION 168

- (Topic 3)

What does Cloud Debugger help an organization do?

- A. Implement code updates in real time without affecting the service level objective (SLO).
- B. Inspect source code in real time without affecting user downtime.
- C. Manage code and accelerate application development.
- D. Analyze live source code during user downtime.

Answer: B

Explanation:

Cloud Debugger is a feature of Google Cloud Platform that lets you inspect the state of an application, at any code location, without stopping or slowing down the running app. Cloud Debugger makes it easier to view the application state without adding logging statements.

NEW QUESTION 171

- (Topic 2)

The customer has applications that do data processing on-premise. They have been built using Ha-doop and Spark. What product should I use on Google Cloud?

- A. Dataproc
- B. Dataflow
- C. Dataprep
- D. Dataplex

Answer: A

Explanation:

Because Dataproc is used to run Hadoop/Spark workloads

NEW QUESTION 173

- (Topic 2)

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

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

Answer: AB

Explanation:

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

NEW QUESTION 175

- (Topic 2)

What issues can arise when organizations integrate third-party systems into their cloud infrastruc-ture?

- A. Third-party systems may not be powerful enough to run many critical business applications.
- B. Without sufficient security measures and regular checks, unsecured third-party systems can pose a threat to data security.
- C. Over-reliance on third-party systems limits an organization's potential for innova-tion.
- D. Third-party systems are less capable of addressing an organization's security re- quirements.

Answer: B

Explanation:

Because unsecured third-party systems are a cybersecurity threat.

NEW QUESTION 177

- (Topic 2)

You are working with a user to set up an application in a new VPC behind a firewall and it is no-ticed that the user is concerned about data egress. Therefore, to provide assistance you want to con-figure the fewest open egress ports. Which of the following statement is correct?

- A. Set up a high-priority (1000) rule that blocks all egress and a low-priority (65534) rule that allows only the appropriate ports.
- B. Set up a low-priority (65534) rule that blocks all egress and a high-priority rule (1000) that allows only the appropriate ports.
- C. Set up a high-priority (1000) rule to allow the appropriate ports.

D. Set up a high-priority (1000) rule that pairs both ingress and egress ports.

Answer: B

Explanation:

Implied rules Every VPC network has two implied firewall rules. These rules exist, but are not shown in the Cloud Console:
 Implied allow egress rule. An egress rule whose action is allow, destination is 0.0.0.0/0, and priority is the lowest possible (65535) lets any instance send traffic to any destination, except for traffic blocked by Google Cloud. A higher priority firewall rule may restrict outbound access. Internet access is allowed if no other firewall rules deny outbound traffic and if the instance has an external IP address or uses a Cloud NAT instance. For more information, see Internet access requirements.

If IPv6 is enabled, the VPC network also has these two implied rules:

- **Implied IPv6 allow egress rule.** An egress rule whose action is allow, destination is ::/0, and priority is the lowest possible (65535) lets any instance send traffic to any destination, except for traffic blocked by Google Cloud. A higher priority firewall rule may restrict outbound access. Internet access is allowed if no other firewall rules deny outbound traffic and if the instance has an external IP address.
- **Implied IPv6 deny ingress rule.** An ingress rule whose action is deny, source is ::/0, and priority is the lowest possible (65535) protects all instances by blocking incoming connections to them. A higher priority rule might allow incoming access.

The implied rules cannot be removed, but they have the lowest possible priorities. You can create rules that override them as long as your rules have higher priorities (priority numbers less than 65535). Because deny rules take precedence over allow rules of the same priority, an ingress allow rule with a priority of 65535 never takes effect.

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

NEW QUESTION 178

- (Topic 2)

A large travel company has thus far invested heavily in their technology team. There is strategic pressure on the company to focus on their core business and innovate to survive in certain geographies and thrive in others. They are evaluating whether a move to Google Cloud will be good for them. Which of these reasons would be relevant for them? (choose two answer)

- A. Application architecture won't be too involved because of serverless options.
- B. The IT team won't have to manage software upgrades, security patches, et
- C. for the VMs.
- D. The IT team won't have to work on procuring and provisioning new hardware and refreshes to existing hardware.
- E. Budgeting won't be an issue since the cloud takes care of billing.

Answer: BC

NEW QUESTION 182

- (Topic 2)

A customer deploys an application to App Engine and needs to check for Open Web Application Security Project (OWASP) vulnerabilities. Which service should be used to accomplish this?

- A. Cloud Armor
- B. Cloud Security Scanner
- C. Binary Authorization
- D. Forseti Security

Answer: B

Explanation:

Web Security Scanner identifies security vulnerabilities in your App Engine, Google Kubernetes Engine (GKE), and Compute Engine web applications. It crawls your application, following all links within the scope of your starting URLs, and attempts to exercise as many user inputs and event handlers as possible. Currently, Web Security Scanner only supports public URLs and IPs that aren't behind a firewall. Web Security Scanner currently supports the App Engine standard environment and App Engine flexible environments, Compute Engine instances, and GKE resources.

Reference link- <https://cloud.google.com/security-command-center/docs/concepts-web-security-scanner-overview>

NEW QUESTION 184

- (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 onl
- C. No additional configuration is require
- 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 189

- (Topic 2)

Considering Different Storage and database options e.g. Cloud Datastore, Cloud SQL, Cloud Storage, etc. Which of the following statements is/are correct? (Select two answer)

- A. Cloud DataStore and Cloud SQL have Terabytes + and Terabytes Capacity respectively.
- B. Cloud Bigtable and Cloud Storage both have Petabytes + capacity.
- C. Cloud Bigtable and Cloud Storage both have not Petabytes + capacity.
- D. None of the above.

Answer: AB

NEW QUESTION 191

- (Topic 2)

DriveSuper Inc. teaches its clients to drive cars and bikes and helps them get their license. They are planning to build a mobile application where users can sign up, plan their schedules, and take stock of progress. They want the onboarding process to be smooth and frictionless, giving users a great experience from the get-go. They want this done as quickly as possible and not be expensive. What is their best option on Google Cloud?

- A. Build the mobile app with Cloud SQL as the backend
- B. Build the mobile app with Cloud Storage as the backend
- C. Build the mobile application with Firebase as the backend
- D. Build the mobile app with Cloud Spanner as the backend

Answer: C

Explanation:

Firebase/Firestore is easy to build and is suitable for user information that could vary in nature.

NEW QUESTION 192

- (Topic 2)

Google offers Firebase, In terms of Firebase Console, any particular message that has to be delivered to a customer at a certain degree of change in behavior can be managed through_____.

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

Answer: B

Explanation:

You can send notification messages using the Notifications composer in the Firebase console. Though this does not provide the same flexibility or scalability as sending messages with the Admin SDK or the HTTP and XMPP protocols, it can be very useful for testing or for highly targeted marketing and user engagement. The Firebase console provides analytics-based A/B testing to help refine and improve marketing messages.

After you have developed logic in your app to receive messages, you can allow non-technical users to send messages per the instructions on the Notifications page in the Firebase Help Center.

NEW QUESTION 194

- (Topic 2)

A financial services company is running an experimental application workload that has a very large number of mathematical calculations involving floating-point numbers. The current application that is running on compute engine is not providing enough speed and throughput. What are the options to increase the processing performance?

- A. Use a serverless option like Cloud Functions that will automatically scale as much as required.
- B. Instead of using a "general purpose" machine family, use "compute-optimized" machine family.
- C. Since processing could also be dependent on reading and writing data to the disk, use a fast Local SSD.
- D. Attach GPUs to the virtual machine for number crunching.

Answer: D

Explanation:

Compute Engine provides graphics processing units (GPUs) that you can add to your virtual machines (VMs). You can use these GPUs to accelerate specific workloads on your VMs such as machine learning and data processing. <https://cloud.google.com/compute/docs/gpus>

NEW QUESTION 196

- (Topic 2)

A developer in your IT team is creating a bucket on Cloud Storage. He is receiving an error that the bucket name already exists. He has checked his project and the few other projects in the organization, The name seems to be entirely unique, What would be the issue?

- A. Bucket names ignore any "." in the name
- B. Look for similar bucket names that have a "." in it.
- C. Previously deleted bucket names in the same project cannot be reused
- D. There must have been an older bucket with the same name.
- E. Bucket names in Cloud storage have to be globally unique
- F. Bucket names are case insensitive- look for bucket names in your org that have a different capitalization.

Answer: C

Explanation:

Bucket names have to be unique across Google Cloud Platform [GCP], including other organizations and projects.

NEW QUESTION 199

- (Topic 2)

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

- 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 covers only Case, phone, and chat support for billing issues only. Reference link- <https://cloud.google.com/support>

NEW QUESTION 202

- (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 203

- (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 204

- (Topic 2)

Which of the following are core components of Anthos?

- A. Infrastructure, container, and cluster management
- B. Secure software supply chain

- C. Multicluster & Configuration management
- D. All of the above are correct.

Answer: D

Explanation:

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

NEW QUESTION 207

- (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 212

- (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 214

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 217

- (Topic 2)

Your client has an on-premises data center. Due to technical limitations, they are unable to scale globally. They have decided to adopt the public cloud. However, they don't want to be locked into any one vendor and, therefore, would like to work with multiple cloud providers. They have used open source container technologies and would like to continue using them.

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

Answer: D

Explanation:

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

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

NEW QUESTION 222

- (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 226

- (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 228

- (Topic 2)

Which of these are defined by the following statement: a contract you have with your end customers, which, if you don't meet, you might even have to pay fines?

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

Answer: A

Explanation:

Service-Level Agreement (SLA)

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

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

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

NEW QUESTION 231

- (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 232

- (Topic 2)

Compute Engine provides machine type recommendations to help you optimize the re- source utilization of your virtual machine (VM) instances. What is this capability called?

- A. App Engine
- B. None of the above
- C. Rightsizing Recommendations
- D. Recommendation Engine

Answer: C

Explanation:

Compute Engine provides machine type recommendations to help you optimize the resource utilization of your virtual machine (VM) instances. These recommendations are generated automatically based on system metrics gathered by the Cloud Monitoring service over the previous 8 days. Use these recommendations to resize your instance's machine type to use the instance's resources more efficiently. This feature is also known as rightsizing recommendations

Reference link- <https://cloud.google.com/compute/docs/instances/apply-machine-type-recommendations-for-instances>

NEW QUESTION 233

- (Topic 2)

Which of the following is/are core storage options available on the Google Cloud Platform?

- A. Cloud Storage and Cloud Data Store
- B. Cloud Spanner
- C. Cloud SQL and Google Big Table
- D. All of the above

Answer: D

Explanation:

Google Cloud Platform has other storage options to meet your needs for structured, unstructured, transactional and relational data. Core storage options: Cloud Storage, Cloud SQL, Cloud Spanner, Cloud Data Store and Google Big Table. Depending on your application, you might want to use one or several of these services to get the job done.

NEW QUESTION 236

- (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 238

- (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 239

- (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 244

- (Topic 2)

You are a program manager in a company you need to submit a bare metal solution order for a secure, high performance connection with a low-latency network fabric. What network information you need to submit the order to Bare Metal Solutions.

- A. IP Ranges for example Client IP Address range used for communication between your Google Cloud and Bare Metal Solution environments.
- B. Google Cloud Project Id that you are using with your bare metal solution environment.
- C. Total number of VLANs you need in your Bare Metal Solution Environment.
- D. All of the above

Answer: D

Explanation:

What Bare Metal Solution provides

Bare Metal Solution is a managed solution that provides purpose-built HPE or Atos bare-metal servers in regional extensions that are connected to Google Cloud by a managed, high-performance connection with a low-latency network fabric.

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 custom, sole-tenancy servers with local SAN, and smart hands support.

The network, which is managed by Google Cloud, includes a low-latency Partner 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 248

- (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. Left 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 252

- (Topic 2)

In terms of Infrastructure as a Service (IaaS) what are the benefits of it?

- A. IaaS offers virtually infinite flexibility and scalability, enterprises can get their work done more efficiently, ensuring faster development life cycles.
- B. IaaS resources are regularly available to businesses when they need the

- C. As a result, enterprises reduce delays when expanding infrastructure and, alternatively, don't waste resources by overbuilding capacity.
- D. IaaS resources are used on demand and enterprises only have to pay for the compute, storage, and networking resources that are actually used, IaaS costs are fairly predictable and can be easily contained and budgeted for.
- E. All of the Above

Answer: D

Explanation:

These are the feature of Infrastructure as a Service (IaaS) It's economical
Because IaaS resources are used on demand and enterprises only have to pay for the compute, storage, and networking resources that are actually used, IaaS costs are fairly predictable and can be easily contained and budgeted for.
It's efficient
IaaS resources are regularly available to businesses when they need them. As a result, enterprises reduce delays when expanding infrastructure and, alternatively, don't waste resources by overbuilding capacity.
It boosts productivity
Because the cloud provider is responsible for setting up and maintaining the underlying physical infrastructure, enterprise IT departments save time and money and can redirect resources to more strategic activities.
It's reliable
IaaS has no single point of failure. Even if any one component of the hardware resources fails, the service will usually still remain available.
It's scalable
One of the biggest advantages of IaaS in cloud computing is the capability to scale the resources up and down rapidly according to the needs of the enterprise. It drives faster time to market
Because IaaS offers virtually infinite flexibility and scalability, enterprises can get their work done more efficiently, ensuring faster development life cycles.

NEW QUESTION 254

- (Topic 2)

You are consulting for a client who is migrating to Google Cloud. They presently have a matrix organization. Their IT environments were managed around projects. Each team had multiple projects. All the projects had a flat structure under the company. What would you advise them when planning for the move?

- A. On Google Cloud, create a folder corresponding to each team
- B. Under that, there could be projects or further sub folders as the team decides.
- C. In terms of not disturbing the project developers and testers, advise them that the strategic decision is to retain the structure on Google Cloud also.
- D. Since a Project could spawn other sub-Projects, on Google Cloud it is better to assign a folder for each Project.
- E. The flat structure is what is currently used in IT organizations, and this can be used as-is which will provide the best results.

Answer: A

Explanation:

Folders for a related group of projects are the recommended approach.
-> A flat structure under the organization node is possible on Google Cloud, but it is not recommended. It becomes tougher to manage.
-> Projects cannot have sub-projects; there can only be resources within Projects.
Reference link- <https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy>

NEW QUESTION 259

- (Topic 2)

Virtual Machine vCPU and memory usage for each of these categories can receive one of the following discounts? (Select Three Answer)

- A. Military Discounts
- B. Spot Instances
- C. Committed-Use
- D. Sustained-Use
- E. Preemptible VMs

Answer: CDE

Explanation:

Sustained, Committed and Preemptible
vCPU and memory usage for each of these categories can receive discounts VM vCPU and memory usage for each of these categories can receive discounts
Sustained-use discounts—Google offers up to 30% off for workloads that run for most of the billing month on GCP services.
Committed-use discounts—users can save up to 57% by committing to use an instance for a certain time period, with no upfront payment and with the flexibility to change instances during the commitment period.
Preemptible VMs—similar to the concept of AWS spot instances, Google offers up to 79% off for Virtual Machines that may be shut down at any time and replaced by others.
Reference link- <https://cloud.google.com/compute/docs/sustained-use-discounts> Reference link- <https://cloud.google.com/compute/docs/instances/signing-up-committed-use-discounts>
Reference link- <https://cloud.google.com/compute/docs/instances/preemptible>

NEW QUESTION 260

- (Topic 2)

You are running a data warehouse on BigQuery. A partner company is offering a recommendation engine based on the data in your data warehouse. The partner company is also running their application on Google Cloud. They manage the resources in their own project, but they need access to the BigQuery dataset in your project. You want to provide the partner company with access to the dataset. What should you do?

- A. Ask the partner to create a Service Account in their project, and have them give the Service Account access to BigQuery in their project.
- B. Create a Service Account in your own project, and grant this Service Account access to BigQuery in your project.
- C. Create a Service Account in your own project, and ask the partner to grant this Service Account access to BigQuery in their project.
- D. Ask the partner to create a Service Account in their project, and grant their Service Account access to the BigQuery dataset in your project.

Answer: D

Explanation:

- if the need is to authenticate the application to access your dataset, it's the application's service account that will be provided during the authentication, so the service account is to be created at their side to run the application

NEW QUESTION 264

- (Topic 2)

A customer has a tens of applications that are dependent on Oracle databases in their on-premise data centers. The customer wants to migrate to Google Cloud. Their long term goal is to move to other cloud native database technologies. What options do they have to initially move their data?

- A. Migrate to a Bare Metal server.
- B. Migrate to Cloud SQL.
- C. Since there is no hosted Oracle solution, leave the Oracle data on-premise while doing analytics on Google Cloud.
- D. Containerize Oracle and run it using Cloud Run.

Answer: B

Explanation:

The Bare Metal solution is the recommended approach. You can deploy Oracle capabilities like clustered databases, replication, and all performance features at licensing costs that are similar to on-premise systems

Choose a Google Cloud bare metal migration strategy

[Send feedback](#)

This article describes the three most common options for migrating your bare metal workloads to Google Cloud along with a framework for understanding your workload requirements. It also explains how to choose the bare metal option that's right for your situation. Finally, it provides practical use cases for each migration strategy.

This article is designed for IT managers and staff who want to understand the capabilities of the Google Cloud offerings [Migrate for Compute Engine](#), [Bare Metal Solution](#), and [Mainframe Modernization](#), and how each can facilitate the migration of bare-metal workloads. The article also discusses an IBM offering for working on Google Cloud.

Migrating to bare metal in Google Cloud serves as an important step toward transforming your IT strategy to focus on the cloud. By running your bare metal workloads closer to Google Cloud services, you can take advantage of those services while implementing your application modernization strategy in parallel.

<https://cloud.google.com/architecture/migrating-bare-metal-workloads>

NEW QUESTION 268

- (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 270

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 272

- (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 276

- (Topic 2)

"With cloud messaging you can Customize and deliver messages accordingly to the predetermined time in the user's local time zone." Comment on the above statement.

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

Answer: D

Explanation:

Firebase Cloud Messaging:

Firebase Cloud Messaging (FCM) is a cross-platform messaging solution that lets you reliably send messages at no cost.

Using FCM, you can notify a client app that new email or other data is available to sync. You can send notification messages to drive user re-engagement and retention. For use cases such as instant messaging, a message can transfer a payload of up to 4000 bytes to a client app.

Key capabilities of Firebase Cloud Messaging:

Send notification messages or data messages: Send notification messages that are displayed to your user. Or send data messages and determine completely what happens in your application code.

Versatile message targeting: Distribute messages to your client app in any of 3 ways—to single devices, to groups of devices, or to devices subscribed to topics.

Send messages from client apps: Send acknowledgments, chats, and other messages from devices back to your server over FCM's reliable and battery-efficient connection channel.

NEW QUESTION 281

- (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 286

- (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 288

- (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 290

- (Topic 2)

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

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

Answer: A

Explanation:

Containers and Kubernetes are ideal for the kind of requirement mentioned here - separate microservices that need to scale independently.

Google Kubernetes Engine (GKE) provides a managed environment for deploying, managing, and scaling your containerized applications using Google infrastructure. The GKE environment consists of multiple machines (specifically, Compute Engine instances) grouped together to form a cluster.

Reference link- <https://cloud.google.com/kubernetes-engine/docs/concepts/kubernetes-engine-overview>

NEW QUESTION 294

- (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 296

- (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 297

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 300

- (Topic 2)

What service is a fully managed real-time messaging service that allows you to send and receive messages between independent applications.

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

Answer: B

Explanation:

Google Cloud Pub/Sub is a scalable, durable event ingestion and delivery system.

-> Pub/Sub allows services to communicate asynchronously, with latencies on the order of 100 milliseconds.

-> Pub/Sub is used for streaming analytics and data integration pipelines to ingest and distribute data. It is equally effective as messaging-oriented middleware for service integration or as a queue to parallelize tasks.

-> Pub/Sub enables you to create systems of event producers and consumers, called publishers and subscribers. Publishers communicate with subscribers asynchronously by broadcasting events, rather than by synchronous remote procedure calls (RPCs).

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

NEW QUESTION 304

- (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 305

- (Topic 2)

An organization wants to evaluate the performance of their entire cloud infrastructure, including metrics like server uptime and response rate reports. Which Google Cloud tool should the organization use?

A. Cloud Trace

B. Cloud Monitoring

C. Cloud Profiler

D. Cloud Debugger

Answer: B

Explanation:

Because Cloud Monitoring enables users to monitor the performance of their entire cloud infrastructure.

NEW QUESTION 310

- (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 311

- (Topic 2)

You are a DevOps Engineer in an E-commerce company that sells products globally, across the countries, Customers buy products, add them to carts or check-in stock from different parts of the world with different timestamps, you need to choose a database that can scale globally without any hassle and lots of developer support, it should be consistent across regions, can scale horizontally to support enormous user, automatically replicates, shards and even auto transaction processing. Which of the following database do you choose?

A. Cloud SQL

B. Cloud Spanner

C. Cloud Firestore.

D. Cloud Storage.

Answer: B

Explanation:

Cloud Spanner:

Fully managed relational database with unlimited scale, strong consistency, and up to 99.999% availability.

- Get all the benefits of relational semantics and SQL with unlimited scale

- Start at any size and scale with no limits as your needs grow

- Enjoy high availability with zero scheduled downtime and online schema changes

- Deliver high-performance transactions with strong consistency across regions and continents

- Focus on innovation, eliminating manual tasks with capabilities like automatic sharding

Automatic sharding

Cloud Spanner optimizes performance by automatically sharding the data based on request load and size of the data. As a result, you can spend less time worrying about how to scale your database and instead focus on scaling your business.

Strong transactional consistency

Purpose-built for external, strong, global transactional consistency.

Regional and multi-regional configurations

No matter where your users may be, apps backed by Cloud Spanner can read and write up-to-date strongly consistent data globally. Additionally, when running a multi-region instance, your database is able to survive a regional failure, and offers industry-leading 99.999% availability.

Online schema changes with no downtime

Cloud Spanner users can make a schema change, whether it's adding a column or adding an index while serving traffic with zero downtime. Hence you now have the flexibility to adapt your database to your business needs without compromising on the availability of your application.

NEW QUESTION 312

- (Topic 2)

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

A. Java Programming Spring Boot Framework for to solve the problem of APIs man-agement.

B. Cloud Functions with Firestore and payment gateways integration development.

C. Apigee API Management

D. Frontend & Backend Development with NodeJs and angular etc.

Answer: C

Explanation:

A top-level idea about Apigee API Management and its offered features can help you solve all questions related to Apigee in Cloud Digital Leader Practice Exam.

Apigee is a platform for developing and managing APIs. By fronting services with a proxy

layer, Apigee provides an abstraction or facade for your backend service APIs and provides security, rate limiting, quotas, analytics, and more.

Apigee services: The APIs that you use to create, manage, and deploy your API proxies. Apigee runtime: A set of containerized runtime services in a Kubernetes cluster that Google maintains. All API traffic passes through and is processed by these services.

NEW QUESTION 316

- (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 319

- (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 are inherent to the public cloud provider

NEW QUESTION 323

- (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 324

- (Topic 1)

Your application is onboarding a number of users. The details of the users vary widely. What kind of database would be most suitable for this use case?

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

Answer: A

Explanation:

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

NEW QUESTION 328

- (Topic 1)

Your organization is migrating to Google Cloud. As part of that effort, it needs to move terabytes of data from on-premises file servers to Cloud Storage. Your organization wants the migration process to be automated and to be managed by Google. Your organization has an existing Dedicated Interconnect connection that it wants to use. Which Google Cloud product or feature should your organization use?

- A. Storage Transfer Service
- B. Migrate for Anthos
- C. BigQuery Data Transfer Service
- D. Transfer Appliance

Answer: A

Explanation:

Reference: <https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets>
 Graphical user interface, text, application, email Description automatically generated
<https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets>

NEW QUESTION 330

- (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 332

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 335

- (Topic 1)

Your organization needs to categorize objects in a large group of static images using machine learning. Which Google Cloud product or service should your organization use?

- A. BigQuery ML
- B. AutoML Video Intelligence
- C. Cloud Vision API
- D. AutoML Tables

Answer: C

Explanation:

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

Derive insights from your images in the cloud or at the edge with AutoML Vision or use pre-trained Vision API models to detect emotion, understand text, and more.

Vision API offers powerful pre-trained machine learning models through REST and RPC APIs. Assign labels to images and quickly classify them into millions of predefined categories. Detect objects and faces, read printed and handwritten text, and build valuable metadata into your image catalog.

NEW QUESTION 340

- (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 341

- (Topic 1)

Your organization consists of many teams. Each team has many Google Cloud projects. Your organization wants to simplify the management of identity and access policies for these projects.

How can you group these projects to meet this goal?

- A. Group each team's projects into a separate domain
- B. Assign labels based on the virtual machines that are part of each team's projects
- C. Use folders to group each team's projects
- D. Group each team's projects into a separate organization node

Answer: C

Explanation:

Folders are nodes in the [Cloud Platform Resource Hierarchy](#). A folder can contain projects, other folders, or a combination of both. Organizations can **use folders to group projects** under the organization node in a hierarchy. For example, your organization might contain multiple departments, each with its own set of Google Cloud resources. Folders allow you to group these resources on a per-department basis. Folders are used to group resources that share common IAM policies. While a folder can contain multiple folders or resources, a given folder or resource can have exactly one parent.

<https://cloud.google.com/resource-manager/docs/creating-managing-folders>

NEW QUESTION 342

- (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 346

- (Topic 1)

What are the network requirements for Private Google Access?

- A. Private Google Access automatically enables any API.
- B. Your network must have appropriate routes for the destination IP ranges used by Google APIs and services.
- C. Both A and B
- D. None of the Above

Answer: B

Explanation:

Network requirements for Private Google Access:

- Because Private Google Access is enabled on a per-subnet basis, you must use a VPC network. Legacy networks are not supported because they don't support subnets.
- Private Google Access does not automatically enable any API. You must separately enable the Google APIs you need to use via the APIs & services page in the Google Cloud Console.
- If you use the private.googleapis.com or therestricted.googleapis.com domain names, you'll need to create DNS records to direct traffic to the IP addresses associated with those domains.
- Your network must have appropriate routes for the destination IP ranges used by Google APIs and services. These routes must use the default internet gateway next hop. If you use the private.googleapis.com or therestricted.googleapis.com domain names, you only need one route (per domain). Otherwise, you'll need to create multiple routes.
- Egress firewalls must permit traffic to the IP address ranges used by Google APIs and services. The implied allow egress firewall rule satisfies this requirement. For other ways to meet the firewall requirement.

NEW QUESTION 349

- (Topic 1)

Your ed-tech start-up was originally launched in a small geography. Any user sign-ups, course progress, tests taken, etc. are captured on a self-managed MySQL database. Every user generates many such transactions. Now you're taking the application globally and preparing for a much larger influx of users from all over the world. The existing MySQL server is unlikely to be able to scale. Which convenient option can be considered?

- A. Migrate to BigQuery
- B. Migrate to Cloud Spanner
- C. Migrate to Cloud SQL
- D. Migrate to Bigtable

Answer: B

Explanation:

Cloud Spanner is a global scale SQL database that scales extremely well. That would be the best choice.

NEW QUESTION 350

- (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 355

- (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 356

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 359

- (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 362

- (Topic 1)

Your organization offers public mobile apps and websites. You want to migrate to a Google Cloud-based solution for checking and maintaining your users' usernames and passwords and controlling their access to different resources based on their identity. Which should your organization choose?

- A. VPN tunnels
- B. Identity Platform
- C. Compute Engine firewall rules
- D. Private Google Access

Answer: B

Explanation:

An identity platform is a modern solution for managing the identities of users and devices in a centralized fashion.

Reference: <https://www.okta.com/blog/2021/07/what-is-an-identity-platform/#:~:text=An%20identity%20platform%20is%20a,%2C%20integrations%2C%20and%20platform%20services>

NEW QUESTION 365

- (Topic 1)

An organization wants to scale their existing virtual machine architecture as quickly as possible. Why should the organization use VMware Engine?

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

Answer: D

Explanation:

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

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

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

Table

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

NEW QUESTION 370

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

Cloud-Digital-Leader Practice Exam Features:

- * Cloud-Digital-Leader Questions and Answers Updated Frequently
- * Cloud-Digital-Leader Practice Questions Verified by Expert Senior Certified Staff
- * Cloud-Digital-Leader Most Realistic Questions that Guarantee you a Pass on Your First Try
- * Cloud-Digital-Leader Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The Cloud-Digital-Leader Practice Test Here](#)