

Exam Questions Cloud-Digital-Leader

Google Cloud Digital Leader exam

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



NEW QUESTION 1

- (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 2

- (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 3

- (Topic 1)

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

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

Answer: C

Explanation:

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

Combining reservations with commitments

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

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

NEW QUESTION 4

- (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] }
```

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 5

- (Topic 1)

Your organization needs to minimize how much it pays for data traffic from the Google network to the internet. What should your organization do?

- A. Choose the Standard network service tier.
- B. Choose the Premium network service tier.

- C. Deploy Cloud VPN.
- D. Deploy Cloud NAT.

Answer: A

Explanation:

Choose the Standard network service tier. While Premium tier is the default for all egress traffic and offers the highest performance, when cost is a consideration. Standard tier is the more economical.

Every cloud deployment needs a network over which to move data. Without a network, you can't view cat videos or upload your selfies, much less allow microservices to talk to one another.

Google Cloud provides a global, scalable, flexible network for your cloud-based workloads and services, and how you utilize that network impacts four critical aspects of your deployment: cost, security, performance and availability.

When designing a reliable, sound, yet cost effective network architecture, you'll want multiple teams within the company to weigh in on these four elements, to determine your priorities. The following tips highlight a few considerations you should think about when architecting your network solution.

<https://cloud.google.com/blog/products/networking/networking-cost-optimization-best-practices>

NEW QUESTION 6

- (Topic 1)

Your manager wants to restrict communication of all virtual machines with internet access; with resources in another network; or with a resource outside Compute Engine. It is expected that different teams will create new folders and projects in the near future.

How would you restrict all virtual machines from having an external IP address?

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

Answer: A

Explanation:

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

NEW QUESTION 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)

Which of the following options is/are correct about Preemptible VMs?

- A. Preemptible VMs don't have fixed pricing.
- B. Both A & B
- C. None of the Above.
- D. You can not use Preemptible VMs at fault-tolerant workloads such as high-performance computing, big data and analytics, continuous integration/continuous delivery (CI/CD), rendering/transcoding, and testing.

Answer: C

Explanation:

Preemptible VMs:

Predictable and low cost

Preemptible VMs are up to 80% cheaper than regular instances. Pricing is fixed so you will always get low cost and financial predictability, without worrying about variable market pricing.

Expand your batch processing

Supplement your regular VMs with lower-cost, preemptible instances to finish your compute-intensive work faster, saving you time and money. Throw preemptible VMs at fault-tolerant workloads such as high performance computing, big data and analytics, continuous integration/continuous delivery (CI/CD), rendering/transcoding, and testing.

Get more from your containers

Containers are naturally stateless and fault tolerant, making them a great fit for preemptible VMs! You save on your containerized workloads today with these affordable compute instances. Take advantage of Google Kubernetes Engine for your containerized workloads and Managed Instance Groups to painlessly and seamlessly recover from preemptions.

Enable it instantly

Simply add --preemptible to the gcloud command line and you're off to the races. There's no bidding to code for, and with per-second billing, just shut down your VMs as soon as you're done.

NEW QUESTION 9

- (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 10

- (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)

Each of the three cloud service models - infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS) - offers benefits between flexibility and levels of management by the cloud provider and the customer.

Why would SaaS be the right choice of service model?

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

Answer: B

Explanation:

Benefits of SaaS

The main benefit of SaaS is that it offloads all infrastructure and application management to the SaaS vendor.

Reference: <https://www.ibm.com/cloud/learn/iaas-paas-saas>

NEW QUESTION 14

- (Topic 1)

What is the difference between Standard and Coldline storage?

- A. Coldline storage is for data for which a slow transfer rate is acceptable.
- B. Standard and Coldline storage have different durability guarantees.
- C. Standard and Coldline storage use different APIs.
- D. Coldline storage is for infrequently accessed data.

Answer: D

Explanation:

Reference: <https://www.msp360.com/resources/blog/google-cloud-nearline-storage-vs-coldline-vs-standard/>

Google Cloud Coldline is a new cold-tier storage for archival data with access frequency of less than once per year. Unlike other cold storage options, Nearline has no delays prior to data access, so now it is the leading solution among competitors.

NEW QUESTION 17

- (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 22

- (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 27

- (Topic 1)

You are a program manager for a team of developers who are building an event-driven application to allow users to follow one another's activities in the app. Each time a user adds himself as a follower of another user, a write occurs in the real-time database.

The developers will develop a lightweight piece of code that can respond to database writes and generate a notification to let the appropriate users know that they have gained new followers. The code should integrate with other cloud services such as Pub/Sub, Firebase, and Cloud APIs to streamline the orchestration process. The application requires a platform that automatically manages underlying infrastructure and scales to zero when there is no activity.

Which primary compute resource should your developers select, given these requirements?

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

Answer: B

Explanation:

Reference: <https://firebase.google.com/docs/functions/use-cases>

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

NEW QUESTION 29

- (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 32

- (Topic 1)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 33

- (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 37

- (Topic 1)

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

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

Answer: B

Explanation:

Create reservations on Compute Engine and BigQuery. You can reserve capacity in advance and use it over a period of time. You could also get a cost advantage.

=> There is no need for involved support. It is self-serve via the console.

=> You can reserve resources in advance when you have the need for it. And when you want to take a pay-per-use approach, that is also possible.

=> It is not a good idea to be lock in/hoard resources; you'll pay unnecessarily for resources. Also, it is difficult to time exactly when the demand will be.

References:

<https://cloud.google.com/compute/docs/instances/reserving-zonal-resources> <https://cloud.google.com/bigquery/docs/reservations-intro>

NEW QUESTION 42

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 44

- (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 47

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 50

- (Topic 1)

Your organization wants to predict the behavior of visitors to its public website. To do that, you have decided to build a machine learning model. Your team has database-related skills but only basic machine learning skills, and would like to use those database skills.

Which Google Cloud product or feature should your organization choose?

- A. BigQuery ML
- B. LookML
- C. TensorFlow
- D. Cloud SQL

Answer: A

Explanation:

Reference: <https://cloud.google.com/architecture/predicting-customer-propensity-to-buy>

NEW QUESTION 51

- (Topic 1)

A retail store has discovered a cost-effective solution for creating self-service kiosks. They can use existing check-out hardware and purchase a virtual customer service application. Why do they also need an API?

- A. To connect the check-out hardware to the public cloud.
- B. To connect the new application with the legacy system.
- C. To migrate all customer data for disaster recovery.
- D. To update the check-out hardware remotely.

Answer: B

Explanation:

APIs can create new business value by connecting legacy systems (the checkout hardware) with new software (the virtual customer service application).

NEW QUESTION 52

- (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 55

- (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 57

- (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 61

- (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 66

- (Topic 1)

Your customer currently has a hybrid cloud setup including their on-premises data center and AWS. They are consolidating all their services on Google Cloud as part of a modernization plan and want to spend less IT effort in the future. There are about 10 MySQL and 25 PostgreSQL databases across the two DCs. What is the best option to for them?

- A. Use the Data Catalog Service to manage the metadata of the databases
- B. Use Cloud Dataflow service and setup Google's Cloud SQL as the sink and the others as the source, which will cause the data to flow in as expected.
- C. Use the Database Migration Service
- D. Use the Bare Metal Solution and copy the databases directly as they are on-premises and on AWS.

Answer: C

Explanation:

Explanation

Database Migration is the right one to use: "Simplifying migrations to Cloud SQL. Now available for MySQL and PostgreSQL migrations, with SQL Server coming soon." Since the customer also doesn't want to manage their own database installations in the future, Cloud SQL is the best option.

Database Migration Service

Simplify migrations to Cloud SQL. Available now for MySQL and PostgreSQL, with SQL Server migrations and Oracle to PostgreSQL migrations in preview.

[Get started](#)[Migration guide](#)

- ✓ Migrate databases to Cloud SQL from on premises, Google Compute Engine, and other clouds
- ✓ Replicate data continuously for minimal downtime migrations
- ✓ Serverless and easy to set up

<https://cloud.google.com/database-migration>

NEW QUESTION 69

- (Topic 1)

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

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

Answer: B**Explanation:**

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

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

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

References:

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

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

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

NEW QUESTION 73

- (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 77

- (Topic 1)

Your organization is developing a mobile app and wants to select a fully featured cloud-based compute platform for it.

Which Google Cloud product or feature should your organization use?

- A. Google Kubernetes Engine
- B. Firebase
- C. Cloud Functions
- D. App Engine

Answer: B

Explanation:

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

Firebase is Google's mobile development platform that empowers you to quickly build and grow your app

NEW QUESTION 80

- (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 85

- (Topic 3)

An organization is looking for a storage solution that will help them serve content to users worldwide. They need a solution that offers a high level of availability. What feature of Cloud Storage would they benefit from?

- A. Global metadata
- B. Object versioning
- C. Data encryption
- D. Multi-regional storage

Answer: D

NEW QUESTION 87

- (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 88

- (Topic 3)

A large retail organization uses traditional technology for their ecommerce website. During peaks in traffic, resources are often underutilized or overprovisioned. They have decided to migrate to cloud technology. What aspect of cloud technology will benefit their ecommerce business?

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

Answer: A

NEW QUESTION 93

- (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 97

- (Topic 3)

An organization needs to search an application's source code to identify a potential issue. The application is distributed across multiple containers. Which Google Cloud product should the organization use?

- A. Google Cloud Console
- B. Cloud Trace
- C. Cloud Monitoring
- D. Cloud Logging

Answer: B

Explanation:

Cloud Trace is supposed to be the correct answer. It's an application performance management tool. It's a Google solution for monitoring application performance. It is a distributed tracing system that helps developers debug or fix and optimize their code

NEW QUESTION 99

- (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 102

- (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 105

- (Topic 3)

An organization wants to build autoscaling web applications without having to manage application infrastructure. Which Google Cloud product should they use?

- A. App Engine
- B. AutoML
- C. Anthos
- D. Apigee

Answer: A

Explanation:

Per Google docs, App Engine, allows for "freeing up your developers with zero server management and zero configuration deployments".
<https://cloud.google.com/appengine>

NEW QUESTION 109

- (Topic 3)

How is service availability measured in the context of cloud technology?

- A. Number of available regions
- B. Percentage of uptime
- C. Speed of response time
- D. Number of downtime incidents

Answer: B

NEW QUESTION 114

- (Topic 3)

An organization's developers are growing increasingly frustrated by the limitations of their on-premises infrastructure. How would they benefit from leveraging cloud technology?

- A. They can expect 100% service availability.
- B. They can avoid the limitations of serverless computing.
- C. They can have new tools to innovate and optimize resource usage.

D. They can optimize maintenance for their on-premises infrastructure.

Answer: C

Explanation:

Google cloud have vast majority of products/tools that you can use to innovate. Additionally, there are products in google that scale automatically based from usage (Ex. App Engine, Cloud Run, etc.)

NEW QUESTION 119

- (Topic 3)

Why is data stored in Google Cloud secure and private?

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

Answer: C

Explanation:

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

NEW QUESTION 124

- (Topic 3)

An organization wants to leverage tooling and automation as part of its new DevOps philosophy. Which operational challenge will this resolve?

- A. Repetitive manual tasks that hinder workflows
- B. Time-consuming supervision of creative tasks
- C. Distribution and supply-chain issues
- D. Defective technical equipment that limits innovation

Answer: A

NEW QUESTION 129

- (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 133

- (Topic 3)

How does Google Cloud ensure that customer data remains secure and private when at rest?

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

Answer: C

Explanation:

Google Cloud commitment to keep the data secure and private:

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

NEW QUESTION 134

- (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 139

- (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 143

- (Topic 3)

An organization relies on online seasonal sales for the majority of their annual revenue. Why should the organization use App Engine for their customer app?

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

Answer: B

NEW QUESTION 145

- (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 147

- (Topic 3)

How does switching from on-premises to the cloud help organizations gain value over time?

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

Answer: A

NEW QUESTION 149

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 150

- (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 153

- (Topic 3)

An organization needs a platform to create custom end-to-end artificial intelligence models. Which Google Cloud product or service should the organization use?

- A. Dataproc
- B. Compute Engine
- C. Recommendations AI
- D. Vertex AI

Answer: D

Explanation:

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

NEW QUESTION 155

- (Topic 3)

An organization is looking for a business intelligence solution that allows individual employees and end users to analyze business data and generate insights. Which Google Cloud product or service should the organization use?

- A. Looker
- B. Cloud Spanner
- C. BigQuery
- D. Dataflow

Answer: A

NEW QUESTION 157

- (Topic 3)

An international bank is looking for a serverless warehouse solution that lets them perform smart analytics. Which Google Cloud product or service should the bank use?

- A. BigQuery
- B. Dataflow
- C. Compute Engine
- D. Cloud Spanner

Answer: A

Explanation:

The international bank should use Google Cloud's BigQuery service, which is a fully managed, serverless data warehouse that allows for high-speed analysis of large datasets. It provides a range of built-in functions for analytics and can easily integrate with other Google Cloud services.

NEW QUESTION 159

- (Topic 3)

An organization wants full control of their virtual machine infrastructure for a custom home-grown application with a product that autoscales and automatically updates. Which Google Cloud product or solution should the organization use?

- A. Cloud Build
- B. Cloud Run
- C. Compute Engine
- D. App Engine

Answer: C

Explanation:

Compute Engine will allow you to have full control of their VM infrastructure and you can autoscale and also apply automatic updates.

NEW QUESTION 162

- (Topic 3)

An organization wants to transform multiple types of structured and unstructured data in the cloud from various sources. The data must be readily accessible for analysis and insights.

Which cloud data storage system should the organization use?

- A. Relational database
- B. Private data center
- C. Data field
- D. Data warehouse

Answer: D

Explanation:

It supports real-time insights. A data warehouse is an enterprise system used for the analysis and reporting of structured and semi-structured data from multiple sources, <https://cloud.google.com/learn/what-is-a-data-warehouse>

NEW QUESTION 167

- (Topic 3)

An organization wants to use all available data to offer predictive suggestions on their website that improve over time. Which method should the organization use?

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

Answer: C

NEW QUESTION 168

- (Topic 3)

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

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

Answer: A

NEW QUESTION 169

- (Topic 3)

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

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

Answer: B

NEW QUESTION 171

- (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 176

- (Topic 3)

An organization wants to use multiple marketing datasets to forecast user acquisition. How should they use cloud technology to gain new insights from the data?

- A. Import the datasets into a custom data warehouse, and then archive old data
- B. Import and selectively archive the datasets in a custom data lake
- C. Separate the datasets and make predictions using machine learning
- D. Combine the datasets and make predictions using machine learning

Answer: D

NEW QUESTION 181

- (Topic 3)

When an organization adopts cloud technology, how does their total cost of ownership (TCO) shift?

- A. Away from cost management toward capital expenditure
- B. Away from operational expenditure toward cost management
- C. Away from capital expenditure toward operational expenditure
- D. Away from operational expenditure toward capital expenditure

Answer: C

NEW QUESTION 182

- (Topic 3)

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

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

Answer: B

NEW QUESTION 184

- (Topic 3)

What is logging within the context of cloud technology?

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

Answer: A

Explanation:

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

NEW QUESTION 188

- (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 190

- (Topic 3)

An organization is searching for an open-source machine learning platform to build and deploy their own custom machine learning applications using TPUs. Which Google Cloud product or service should the organization use?

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

Answer: A

Explanation:

<https://en.wikipedia.org/wiki/TensorFlow> TensorFlow is a free and open-source software library for machine learning and artificial intelligence. Developer Google Brain Team

NEW QUESTION 195

- (Topic 3)

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

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

Answer: A

NEW QUESTION 200

- (Topic 3)

An organization wants to digitize and share large volumes of historical text and images. Why is a public cloud a better option than an on-premises solution?

- A. In-house hardware management
- B. Provides physical encryption key
- C. Cost-effective at scale
- D. Optimizes capital expenditure

Answer: C

NEW QUESTION 204

- (Topic 3)

A retail organization has moved all of their inventory data to a relational database in the cloud.

What functionality does a relational database offer?

- A. It analyzes unstructured data which can then be accessed in multiple regions
- B. It stores transactional data which can then be accessed electronically
- C. It stores large amounts of raw data in its original format
- D. It rapidly analyzes large and multi-dimensional datasets

Answer: B

Explanation:

A relational database offers the functionality of storing transactional data, which can then be accessed electronically. Relational databases store structured data that can be organized in tables with defined relationships between them. This makes them well-suited for transactional data, such as inventory data, that needs to be accessed and updated frequently.

NEW QUESTION 207

- (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 211

- (Topic 3)

What is artificial intelligence?

- A. Any system that ingests data in real time
- B. Any system that automatically structures data
- C. Any system capable of a task that requires smart analytics to generate predictions
- D. Any system capable of a task that normally requires human cognition

Answer: D

NEW QUESTION 216

- (Topic 3)

How does Cloud SQL help organizations create business insights?

- A. Integrates with business intelligence and analytics platforms
- B. Generates predictions using machine learning models
- C. Generates real-time charts and intelligent analytics
- D. Transforms business data from unstructured to structured

Answer: A

Explanation:

<https://cloud.google.com/sql/docs/postgres/using-query-insights>

NEW QUESTION 220

- (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 221

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 224

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 227

- (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 231

- (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 232

- (Topic 2)

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

- 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 innovation.
- D. Third-party systems are less capable of addressing an organization's security requirements.

Answer: B

Explanation:

Because unsecured third-party systems are a cybersecurity threat.

NEW QUESTION 235

- (Topic 2)

Your client is building a custom machine learning pipeline to identify lesions in the lungs based on x-rays. Different teams of data scientists are sharing common source data and building many ver-sions of ML models. Which of these Cloud Storage options would be best for them?

- A. Retain the data in use in a single region bucket with nearline storag
- B. Retain the data in use in a dual-region bucket.
- C. Retain the data in use in a single region bucket with standard storage.
- D. Retain the data in use in a multi-region bucket.
- E. Retain the data in use in a dual-region bucket.

Answer: B

Explanation:

Integrated repository for analytics and ML: The highest level of availability and performance within a single region is ideal for compute, analytics, and machine learning workloads in a particular region. Cloud Storage is also strongly consistent, giving you confidence and accuracy in analytics workloads.

Standard storage

Standard storage is best for data that is frequently accessed ("hot" data) and/or stored for only brief periods of time.

When used in a region, Standard storage is appropriate for storing data in the same location as [Google Kubernetes Engine clusters](#) or [Compute Engine instances](#) that use the data. Co-locating your resources maximizes the performance for data-intensive computations and can reduce network charges.

When used in a dual-region, you still get optimized performance when accessing Google Cloud products that are located in one of the associated regions, but you also get the improved availability that comes from storing data in geographically separate locations.

When used in a multi-region, Standard storage is appropriate for storing data that is accessed around the world, such as serving website content, streaming videos, executing interactive workloads, or serving data supporting mobile and gaming applications.

Availability

The availability of Standard storage data is:

Location Type	Availability SLA ¹	Typical monthly availability
multi-region	99.95%	>99.99%
dual-region	99.95%	>99.99%

Table Description automatically generated with medium confidence <https://cloud.google.com/storage/docs/storage-classes>

NEW QUESTION 237

- (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 geogra-phies 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 238

- (Topic 2)

App Engine has been deployed in your customers GCP cloud deployment. The customer would like to know more about the benefits of App Engine Flexible. Please advise them on the benefits of App Engine Flexible (Select Two Answers)

- A. Supports autoscaling
- B. Supports Docker containers
- C. Supports mainframe connectivity
- D. Source code is written in specific versions of the supported programming lan-guages only

Answer: AB

Explanation:

Autoscaling is supported in both Flexible and Standard environments. Flexible Environment does run a Docker container that includes a custom runtime or source code written in other programming languages.

Reference link - <https://cloud.google.com/appengine/docs/the-appengine-environments>

NEW QUESTION 241

- (Topic 2)

You are working with a government agency. A web application serves users of the country. It al-lows citizens to receive certain services in providing their national identity. Citizens have com-plained that they are seeing delays in web page loading compared to before. On investigating, they are seeing a lot of spurious traffic

coming in from a few IPs which they have identified as for-eign. What should they do?

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

Answer: B

Explanation:

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

NEW QUESTION 244

- (Topic 2)

Your company has made plans to roll out OpenShift, a Kubernetes platform solution offered by IBM Red Hat, across all its on-premises and public cloud environments. Given that you are the lead architect responsible for your company's GCP deployments, what type of shared responsibility model will this deployment entail for you?

- A. SaaS
- B. On premises
- C. PaaS
- D. IaaS

Answer: D

Explanation:

The key to remember here is that for a service provided (GCP in this case) to take responsibility for its PaaS, it must offer the service as a managed service. GCP offers its own Kubernetes platform called GKE. But OpenShift is not a Google-offered PaaS solution. As such, Google will not take responsibility for the back-end operations and design of your OpenShift environments. You will need to manage all the VMs that OpenShift will provision as part of its GCP deployment. So this is an IaaS deployment from a shared responsibility model perspective.

NEW QUESTION 246

- (Topic 2)

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

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

Answer: B

Explanation:

Logged information

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

Available audit logs

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

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

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

For fuller descriptions of the audit log types, see [Types of audit logs](#).

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

NEW QUESTION 250

- (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 254

- (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 259

- (Topic 2)

Which Firebase quality tools help the developer track, prioritize & fix stability issues that erode the application quality?

- A. Performance
- B. App Distribution
- C. Crashlytics
- D. Test Lab

Answer: C

Explanation:

Firebase Crashlytics:

Get clear, actionable insight into app issues with this powerful crash reporting solution for iOS, Android, and Unity.

Firebase Crashlytics is a lightweight, real-time crash reporter that helps you track, prioritize, and fix stability issues that erode your app quality. Crashlytics saves you troubleshooting time by intelligently grouping crashes and highlighting the circumstances that lead up to them.

Find out if a particular crash is impacting a lot of users. Get alerts when an issue suddenly increases in severity. Figure out which lines of code are causing crashes.

NEW QUESTION 262

- (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 266

- (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 267

- (Topic 2)

certain devices for cracks, rust, etc. Some of these issues are difficult to identify for a human and your company has seen increasing customer complaints - the customer has paid for an inspection and the field agent said there was no problem, but it later turned out there actually was. The team has come up with a

proposal to engage AI to identify issues. On evaluating the existing system, it is seen that the mobile phone network connection is not good or consistent. What solution can work for them?

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

Answer: A

Explanation:

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

NEW QUESTION 270

- (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 clou
- 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 271

- (Topic 2)

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

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

Answer: A

Explanation:

Service-Level Agreement (SLA)

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

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

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

NEW QUESTION 273

- (Topic 2)

Which of the following statements describe the features of a preemptible VM in-stance? (Select Three Answer)

- A. Instance is alive for no more than 12 hours
- B. Can be pre-empted with a 30 minute notice
- C. Can be pre-empted with a 30 second notice
- D. Discounted Significantly
- E. Instance is alive for no more than 24 hours
- F. Can use free tier credits

Answer: CDE

Explanation:

Instance is alive for no more than 24 hours, Can be pre-empted with a 30 second notice, Discounted Significantly. Preemptible VM is an instance that you can create and run at a lower cost than normal instances. However, Compute Engine might stop (pre-empt) these instances if it requires access to those resources for other tasks. Preemptible instances are excess Compute Engine capacity, so their availability varies with usage. Live at most 24 hours Can be pre-empted with a 30 second notification via API and are Discounted significantly Reference link- <https://cloud.google.com/compute/docs/instances/preemptible>

NEW QUESTION 274

- (Topic 2)

A customer in the European Union region is very clear that their data should not go outside the European Union. Their end users are spread all over the European U. They have to choose a storage option that serves all the users within Asia via web browsers as quickly as possible. Which storage option will work for them?

Multi-regions

Multi-Region Name	Multi-Region Description
ASIA	Data centers in Asia
EU	Data centers within member states of the European Union*
US	Data centers in the United States

- A. Cloud Storage with a single region that is known to be within the European U
- B. Cloud Filestore is connected to virtual machines which are guaranteed to be within the European U
- C. Cloud Storage with the multi-region option of European U
- D. Cloud Storage with the dual-region option of European U

Answer: C

Explanation:

Multi-region option will use multiple datacenters that are within the European Union. More regions will also help with lower latency since users are spread across the European U.
<https://cloud.google.com/storage/docs/locations#considerations>

NEW QUESTION 275

- (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 276

- (Topic 2)

Which of the following is / are true for Preemptible Instances.

- A. Preemptible Instances have no Service Level Agreement (Compute Engine SLA).
- B. Google Cloud Free Tier credits for compute engine do not apply to preemptible in-stances.
- C. Preemptible instances can't live migrate to a regular VM instance, or be set to automatically restart when there is a maintenance event.
- D. All of the above.

Answer: D

Explanation:

Preemptible instances function like normal instances but have the following limitations:
 -> Compute Engine might stop preemptible instances at any time due to system events. The probability that Compute Engine will stop a preemptible instance for a system event is generally low, but might vary from day to day and from zone to zone depending on current conditions.
 -> Compute Engine always stops preemptible instances after they run for 24 hours. Certain actions reset this 24-hour counter.
 -> Preemptible instances are finite Compute Engine resources, so they might not always be available.
 -> Preemptible instances can't live migrate to a regular VM instance, or be set to automatically restart when there is a maintenance event.

-> Due to the above limitations, preemptible instances are not covered by any Service Level Agreement (and, for clarity, are excluded from the Compute Engine SLA).
-> The Google Cloud Free Tier credits for Compute Engine do not apply to preemptible instances.

NEW QUESTION 277

- (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 279

- (Topic 2)

Considering Google Cloud Storage different Options which of the following is / are correct on the basis of their real world use cases?

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

Answer: D

Explanation:

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

NEW QUESTION 282

- (Topic 2)

You are a program manager in a company you need to submit a bare metal solution order for a se-cure, 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 environ-ment.
- 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 285

- (Topic 2)

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

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

Answer: B

Explanation:

Preemptible VMs have all these features

Simple configuration

Create a preemptible instance simply by flipping a bit via command, API, or developer console.

Easy extensibility

Attach GPUs and local SSDs to preemptible instances for additional performance and savings.

Graceful shutdown

Compute Engine gives you 30 seconds to shut down when you're preempted, letting you save your work in progress for later.

Large scale computing

Spin up as many instances as you need and turn them off when you're done. You only pay for what you use.

Quickly reclaim capacity

Managed instance groups automatically recreate your instances when they're preempted (if capacity is available).

Fixed pricing

Preemptible VMs have fixed pricing up to 80% off regular instances. They show up on your bill separately so you'll see just how much you're saving.

NEW QUESTION 290

- (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 291

- (Topic 2)

A client is currently running software on their on-premise systems that is bound by a certain

type of license. They are allowed to run the software on virtualized machines. However, they cannot run them on virtualized machines that are shared by two different companies, teams, or projects. What option do they have on Google Cloud?

A. Google Cloud is a public cloud accessed by multiple customers.

B. Allocate a Bare Metal machine.

C. Setup exclusive login to the VM with self-generated security keys.

D. Allocate sole-tenant nodes

Answer: D

Explanation:

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

<https://cloud.google.com/compute/docs/nodes/sole-tenant-nodes>

NEW QUESTION 292

- (Topic 2)

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

A. Since the app is anyways going to be public, create permanent credentials for the UX testers that they can conveniently use each time.

B. It would be a security issue to have users come and go

C. Recommend that the test-ers be permanently hired to plug the vulnerability issue.

D. It would be a security issue to have users come and go

E. Recommend that the test-ers be permanently hired to plug the vulnerability issue.

F. Create a Group with the permissions required to do the test and record their in-put

G. When users arrive each week, add them to the group and after the testing period, remove them from the group.

Answer: D

Explanation:

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

NEW QUESTION 295

- (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 298

- (Topic 2)

Which of the following statements is/are true about Cloud Spanner offered by Google Cloud Platform.

- A. It can scale horizontally to support additional capacity.
- B. It comes with Zero Downtime, No Maintenance windows, and is proven for large and small workloads.
- C. You don't need to shard or replicate data.
- D. All of the above.

Answer: D

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.

NEW QUESTION 300

- (Topic 2)

In terms of Cloud SQL for MySQL Features offered by Google Cloud Platform which of the statements is/are correct?

- A. Do not support Private IP (private service access).
- B. Customer data is encrypted on Google's internal networks and in database tables, temporary files, and backups.
- C. Do not Provide automated and on-demand backups and point-in-time recovery.
- D. None of the above

Answer: B

Explanation:

Cloud SQL for MySQL:

Features

- Fully managed MySQL Community Edition databases in the cloud.
- Cloud SQL instances support MySQL 8.0, 5.7 (default), and 5.6, and provide up to 624 GB of RAM and 64 TB of data storage, with the option to automatically increase the storage size, as needed.
- Create and manage instances in the Google Cloud Console.
- Instances are available in the Americas, EU, Asia, and Australia.
- Customer data is encrypted on Google's internal networks and in database tables, temporary files, and backups.
- Support for secure external connections with the Cloud SQL Auth proxy or with the SSL/TLS protocol.
- Support for private IP (private services access).
- Data replication between multiple zones with automatic failover.
- Import and export databases using mysqldump, or import and export CSV files.
- Support for MySQL wire protocol and standard MySQL connectors.
- Automated and on-demand backups and point-in-time recovery.
- Instance cloning.
- Integration with Google Cloud's operations suite logging and monitoring.

NEW QUESTION 301

- (Topic 2)

An e-commerce company's business has been booming. To keep up with the growth the IT team also grew. Many new people are being added and new resources are being set up. The CIO is in conversation with you over coffee one day and expresses her growing concern that they might be moving too fast. Their security checks and policies have not kept pace. She worries that somebody would make a misconfiguration or compliance violation thus exposing the company to data and privacy loss. What can you advise her?

- A. Use Cloud Identity-Aware Proxy to allow only specific users to access the data.
- B. Use Security Command Center to have a centralized view of assets and get notified on misconfigurations and vulnerabilities.
- C. Use Cloud Data Loss Prevention to prevent the loss of any data.
- D. Use Cloud Armor to block any DDoS attacks that could be a threat.

Answer: B

Explanation:

Security Command Center is the right tool for this use case. It can check resources for security issues and notify you when issues are found.
<https://cloud.google.com/security-command-center>

NEW QUESTION 304

- (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 309

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 312

- (Topic 2)

A customer has an application running in virtual machines. They are migrating this application to Google Cloud. They have previously had scaling issues when on-premises as VMs had to be pre-allocated. Capacity planning was repeatedly off mark - it's either too many VMs or too less. They want to match the capacity to demand while keeping the application running always. They don't have the time or budget to re-architect the systems using containers and Kubernetes at the moment. What would be your recommendation?

- A. Run a load test on Compute Engine VM
- B. Get an estimate of usage
- C. Then plan for a VM capacity of 25% above the load test value.
- D. Use the Managed Instance Group with Compute Engine
- E. Inform them that new-age companies are using microservices, containers, and Kubernetes for this and they can plan to rewrite the app quickly.
- F. Inform them that using a serverless option will take care of the scaling and they can move to Cloud Run or App Engine.

Answer: B

Explanation:

Scalability. When your apps require additional compute resources, autoscaled MIGs can automatically grow the number of instances in the group to meet demand. If demand drops, autoscaled MIGs can automatically shrink to reduce your costs

Instance groups

[Send feedback](#)

An instance group is a collection of virtual machine (VM) instances that you can manage as a single entity.

Compute Engine offers two kinds of VM instance groups, managed and unmanaged:

- **Managed instance groups (MIGs)** let you operate apps on multiple identical VMs. You can make your workloads scalable and highly available by taking advantage of automated MIG services, including: autoscaling, autohealing, regional (multiple zone) deployment, and automatic updating.
- **Unmanaged instance groups** let you load balance across a fleet of VMs that you manage yourself.

[Return to top](#)

<https://cloud.google.com/compute/docs/instance-groups>

NEW QUESTION 316

- (Topic 2)

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

these applications would be suitable for that?

- A. Use Cloud Vision API - write code to identify the text blocks, copy the data, and store it
- B. Use TensorFlow - write code that will identify the type of visa and the bounding text block
- C. Copy the data and then store it.
- D. Use AutoML - upload other images of visas and run the model creation process which will automatically identify the visas
- E. Use Data Labeling service - outsource the work of marking and extracting the information to others.

Answer: A

Explanation:

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

NEW QUESTION 318

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 320

- (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 321

- (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 322

- (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 324

- (Topic 2)

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

You are working in a company where you need to store Terabytes of Image Data daily and process them e.g. Taking photos of the entire planet 24 hours every day with satellite and

sending data to data centres to store and process it. Which of the following would be the best combination for your infrastructure.

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

Answer: C

Explanation:

The above is a real world example of a company named Planet, where they sent around 80+ satellites to take pictures of earth every day, 24 hours. They run around 40,000 preemptible VMs concurrently.

Preemptible instances function like normal instances but have the following limitations: Compute Engine might stop preemptible instances at any time due to system events. The probability that Compute Engine will stop a preemptible instance for a system event is generally low, but might vary from day to day and from zone to zone depending on current conditions.

Compute Engine always stops preemptible instances after they run for 24 hours. Certain actions reset this 24-hour counter.

Preemptible instances are finite Compute Engine resources, so they might not always be available.

Preemptible instances can't live migrate to a regular VM instance, or be set to automatically restart when there is a maintenance event.

Due to the above limitations, preemptible instances are not covered by any Service Level Agreement (and, for clarity, are excluded from the Compute Engine SLA).

The Google Cloud Free Tier credits for Compute Engine do not apply to preemptible instances.

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

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

NEW QUESTION 325

- (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 326

- (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 330

- (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 331

- (Topic 1)

The operating systems of some of your organization's virtual machines may have a security vulnerability.

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

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

Answer: A

Explanation:

Security Health Analytics and Web Security Scanner detectors generate vulnerabilities findings that are available in Security Command Center. Your ability to view and edit findings is determined by the Identity and Access Management (IAM) roles and permissions you are assigned. For more information about IAM roles in Security Command Center.

Reference link:-

<https://cloud.google.com/security-command-center/docs/concepts-vulnerabilities-findings>

NEW QUESTION 333

- (Topic 1)

Your organization runs a distributed application in the Compute Engine virtual machines. Your organization needs redundancy, but it also needs extremely fast communication (less than 10 milliseconds) between the parts of the application in different virtual machines.

Where should your organization locate this virtual machines?

- A. In a single zone within a single region
- B. In different zones within a single region
- C. In multiple regions, using one zone per region
- D. In multiple regions, using multiple zones per region

Answer: B

Explanation:

Multi zone is also redundant within the region and it provides the lowest latency.

Reference link:-

<https://cloud.google.com/solutions/best-practices-compute-engine-region-selection>

NEW QUESTION 338

- (Topic 1)

Your organization wants to migrate your on-premises environment to Google Cloud. The on-premises environment consists of containers and virtual machine instances. Which Google Cloud products can help to migrate the container images and the virtual machine disks?

- A. Compute Engine and Filestore
- B. Artifact Registry and Cloud Storage
- C. Dataflow and BigQuery
- D. Pub/Sub and Cloud Storage

Answer: A

Explanation:

Reference: <https://cloud.google.com/compute/docs/import/importing-virtual-disks>
Graphical user interface, text, application, email Description automatically generated

NEW QUESTION 339

- (Topic 1)

Your organization runs an application on virtual machines in Google Cloud. This application processes incoming images. This activity takes hours to create a result for each image. The workload for this application normally stays at a certain baseline level, but at regular intervals it spikes to a much greater workload. Your organization needs to control the cost to run this application.

What should your organization do?

- A. Purchase committed use discounts for the baseline load
- B. Purchase committed use discounts for the expected spike load
- C. Leverage sustained use discounts for your virtual machines
- D. Run the workload on preemptible VM instances

Answer: C

Explanation:

The idea of the Sustained Use discount is that the longer you run a VM instance in any given month, the bigger discount you will get from the list price.

Reference: <https://www.parkmycloud.com/blog/google-sustained-use-discounts/>

NEW QUESTION 344

- (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 348

- (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 350

- (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 355

- (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 356

- (Topic 1)

An organization wants to search for and share plug-and-play AI components which can easily build ML services into their project. Which Google Cloud product should the organization use?

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

Answer: B

Explanation:

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

NEW QUESTION 359

- (Topic 1)

Your company's development team is building an application that will be deployed on Cloud Run. You are designing a CI/CD pipeline so that any new version of the application can be deployed in the fewest number of steps possible using the CI/CD pipeline you are designing. You need to select a storage location for the images of the application after the CI part of your pipeline has built them.

What should you do?

- A. Create a Compute Engine image containing the application
- B. Store the images in Container Registry
- C. Store the images in Cloud Storage
- D. Create a Compute Engine disk containing the application

Answer: B

Explanation:

Reference: <https://cloud.google.com/container-registry/docs/pushing-and-pulling>

NEW QUESTION 364

- (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 368

- (Topic 1)

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

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

Answer: A

Explanation:

Reference: <https://cloud.google.com/security/infrastructure/design>

The Teams need to be segmented to have visibility on the resources each team consumes

NEW QUESTION 371

- (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 372

- (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 375

- (Topic 1)

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

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

Answer: B

Explanation:

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

Introduction to the Organization Policy Service

[Send feedback](#)

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

Benefits

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

References link:

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

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

NEW QUESTION 376

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 377

- (Topic 1)

A prospect wants to be able to store and analyze data. Their analysts already know SQL, but are not familiar with other technologies. Which of these databases can the analysts use without additional training?

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

Answer: B

Explanation:

Spanner, Cloud SQL, BigQuery

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

NEW QUESTION 380

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 382

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 387

- (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 392

- (Topic 1)

Your organization is developing an application that will manage payments and online bank accounts located around the world. The most critical requirement for your database is that each transaction is handled consistently. Your organization anticipates almost unlimited growth in the amount of data stored. Which Google Cloud product should your organization choose?

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

Answer: D

Explanation:

Features of Cloud Spanner

Reference: <https://k21academy.com/google-cloud/cloud-sql-vs-cloud-spanner/>

NEW QUESTION 397

- (Topic 1)

How do Migrate for Compute Engine and Migrate for Anthos differ?

- A. Unlike Migrate for Anthos, Migrate for Compute Engine assumes that the migration source is VMware vSphere.
- B. Migrate for Compute Engine charges for ingress, but Migrate for Anthos does not.
- C. Migrate for Compute Engine is closed source, and Migrate for Anthos is open source.
- D. Migrate for Anthos migrates to containers, and Migrate for Compute Engine migrates to virtual machines.

Answer: D

Explanation:

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

Migrate workloads to Compute Engine with Migrate for Compute Engine. Migrate from Compute Engine to containers with Migrate for Anthos and GKE.

This method makes sense, for instance, in cases where you want to conduct a data-center migration and migrate all workloads into Compute Engine, and only at a second stage selectively modernize suitable workloads to containers.

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

<https://cloud.google.com/migrate/containers/docs/architecture>

NEW QUESTION 398

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

Cloud-Digital-Leader Practice Exam Features:

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