

Google

Exam Questions Cloud-Digital-Leader

Google Cloud Digital Leader exam



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

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 3

- (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/999999',
    'displayName': 'new-folder'
  }
}

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

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

Rectangular Snip

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

Reference link- <https://stackoverflow.com/questions/59460623/how-to-create-a-folder-a-project-under-it-with-deployment-manager-google-cloud>

NEW QUESTION 4

- (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 5

- (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 6

- (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 7

- (Topic 1)

Which of the following is/are true about Anthos?

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

Answer: D

Explanation:

Anthos :

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

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

Key features:

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

NEW QUESTION 8

- (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 9

- (Topic 1)

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

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

Answer: A

Explanation:

Explanation

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

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

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

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

=> Cloud Functions- Cloud Functions is a useful serverless endpoint. However, Pub/Sub is better in this case because it can also retain messages for a set period

if it was not possible to deliver it first time.

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

NEW QUESTION 10

- (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 10

- (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 11

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 13

- (Topic 1)

Your organization is developing a plan for migrating to Google Cloud.

What is a best practice when initially configuring your Google Cloud environment?

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

Answer: B

Explanation:

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

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

NEW QUESTION 15

- (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 19

- (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 22

- (Topic 1)

Your organization is moving an application to Google Cloud. As part of that effort, it needs to migrate the application's working database from another cloud provider to Cloud SQL. The database runs on the MySQL engine. The migration must cause minimal disruption to users. Data must be secured while in transit. Which should your organization use?

- A. BigQuery Data Transfer Service
- B. MySQL batch insert
- C. Database Migration Service
- D. Cloud Composer

Answer: C

Explanation:

Reference: <https://aws.amazon.com/dms/>

NEW QUESTION 23

- (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 27

- (Topic 1)

As your organization increases its release velocity, the VM-based application upgrades take a long time to perform rolling updates due to OS boot times. You need to make the application deployments faster.

What should your organization do?

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

Answer: B

NEW QUESTION 32

- (Topic 1)

Your organization wants to run a container-based application on Google Cloud. This application is expected to increase in complexity. You have a security need for fine-grained control of traffic between the containers. You also have an operational need to exercise fine-grained control over the application's scaling policies. What Google Cloud product or feature should your organization use?

- A. Google Kubernetes Engine cluster
- B. App Engine
- C. Cloud Run
- D. Compute Engine virtual machines

Answer: A

Explanation:

Google Kubernetes Engine GKE seems a better fit since the requirement is for "security need for fine-grained control of traffic between the containers" and "fine-grained control over scaling policies". Such level of control is easier on GKE than Cloud Run.

When it comes to managed Kubernetes services, Google Kubernetes Engine (GKE) is a great choice if you are looking for a **container orchestration platform** that offers advanced scalability and configuration flexibility. GKE gives you complete control over every aspect of container orchestration, from networking, to storage, to how you set up observability—in addition to supporting stateful application use cases. However, if your application does not need that level of cluster configuration and monitoring, then fully managed **Cloud Run** might be the right solution for you.

Fully managed Cloud Run is an ideal **serverless platform** for stateless containerized microservices that don't require Kubernetes features like namespaces, co-location of containers in pods (sidecars) or node allocation and management.

Reference link- <https://cloud.google.com/blog/products/containers-kubernetes/when-to-use-google-kubernetes-engine-vs-cloud-run-for-containers>

NEW QUESTION 33

- (Topic 1)

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

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

Answer: C

Explanation:

Graphical user interface, text, application, Word, email

Direct Peering overview

[Send feedback](#)

Direct Peering enables you to establish a direct **peering**  connection between your business network and Google's edge network and exchange high-throughput cloud traffic.

This capability is available at any of more than 100 locations in 33 countries around the world. For more information about Google's edge locations, see [Google's peering site](#).

When established, Direct Peering provides a direct path from your on-premises network to Google services, including Google Cloud products that can be exposed through one or more public IP addresses. Traffic from Google's network to your on-premises network also takes that direct path, including traffic from VPC networks in your projects. Google Cloud customers must request that direct egress pricing be enabled for each of their projects after they have established Direct Peering with Google. For more information, see [Pricing](#).

Direct Peering exists outside of Google Cloud. Unless you need to access Google Workspace applications, the recommended methods of access to Google Cloud are [Dedicated interconnect](#) or [Partner interconnect](#).

For a description of the differences between Direct Peering and Cloud Interconnect, see the [comparison table](#).

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

NEW QUESTION 35

- (Topic 1)

What are the key features of Google Cloud Identity?

- A. Multi-factor authentication (MFA)
- B. Single sign-on (SSO)

- C. Works with your favorite apps and Endpoint management
- D. All of the Above

Answer: D

Explanation:

Cloud Identity:

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

- Give users easy access to apps with single sign-on.
- Multi-factor authentication protects user and company data.
- Endpoint management enforces policies for personal and corporate devices

KEY FEATURES :

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

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

Endpoint management

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

Single sign-on (SSO)

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

Works with your favorite apps

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

NEW QUESTION 37

- (Topic 1)

Your organization runs many workloads in different Google Cloud projects, each linked to the same billing account. Each project's workload costs can vary from month to month, but the overall combined cost of all projects is relatively stable. Your organization needs to optimize its cost.

What should your organization do?

- A. Purchase a commitment per project for each project's usual minimum
- B. Create a billing account per project, and link each project to a different billing account
- C. Turn on committed use discount sharing, and create a commitment for the combined usage
- D. Move all workloads from all different projects into one single consolidated project

Answer: C

Explanation:

Turn on committed use discount sharing, and create a commitment for the combined usage

Sharing your committed use discounts across all your projects reduces the overhead of managing discounts on a per-project basis, and maximizes your savings by pooling all your discounts across your projects' resource usage. If you have multiple projects that share the same Cloud Billing account, you can enable committed use discount sharing so all of your projects within that Cloud Billing account share all of your committed use discount contracts. Your sustained use discounts are also pooled at the same time. That is, sustained use discounts are calculated using the total resources across these projects, rather than just the resources within a single project.

Sharing committed use discounts across projects

Sharing your committed use discounts across all your projects reduces the overhead of managing discounts on a per-project basis, and maximizes your savings by pooling all your discounts across your projects' resource usage.

If you have multiple projects that share the same Cloud Billing account, you can [enable committed use discount sharing](#) so all of your projects within that Cloud Billing account share all of your committed use discount contracts. Your sustained use discounts are also pooled at the same time. That is, sustained use discounts are calculated using the total resources across these projects, rather than just the resources within a single project.

For example, if you purchase two commitment contracts for a total of 160 cores, and you run 200 cores during the month, you will receive committed use discounts for 160 cores across the projects that used them. The additional 40 cores will be billed at on-demand, non-committed use rates. After you purchase a set amount of commitments, you're billed for those commitments monthly, even if you don't use them. For example, if you purchase commitments for 160 cores, you're billed the committed use rates for those 160 cores for the whole month, even if don't use them. See [Understanding discount sharing](#) for cost-saving utilization recommendations.

Reference link- https://cloud.google.com/compute/docs/instances/signing-up-committed-use-discounts#sharing_committed_use_discounts_across_projects

NEW QUESTION 42

- (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 47

- (Topic 1)

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

What should your organization do?

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

Answer: A

Explanation:

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

NEW QUESTION 49

- (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 52

- (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 53

- (Topic 1)

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

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

Answer: B

Explanation:

Apply the Principle of Least Privilege and only provide read permissions on only the required buckets. No more, no less
<https://cloud.google.com/storage/docs/access-control/iam-roles>

NEW QUESTION 57

- (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 60

- (Topic 1)

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

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

Answer: A

Explanation:

BigQuery ML lets you create and execute machine learning models in BigQuery using standard SQL queries.
 Reference: <https://cloud.google.com/bigquery-ml/docs/introduction#:~:text=BigQuery%20ML%20lets%20you%20create,the%20need%20to%20move%20data>

Graphical user interface, text, application, email Description automatically generated
<https://cloud.google.com/bigquery-ml/docs/introduction>

NEW QUESTION 65

- (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 69

- (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 71

- (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 73

- (Topic 3)

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

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

Answer: C

NEW QUESTION 74

- (Topic 3)

An organization has servers running mission-critical workloads on-premises around the world. They want to modernize their infrastructure with a multi-cloud architecture.

What benefit could the organization experience?

- A. Ability to disable regional network connectivity during cyber attacks
- B. Ability to keep backups of their data on-premises in case of failure
- C. Full management access to their regional infrastructure
- D. Reduced likelihood of system failure during high demand events

Answer: D

NEW QUESTION 75

- (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 77

- (Topic 3)

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

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

Answer: B

NEW QUESTION 79

- (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 83

- (Topic 3)

An organization wants to use BigQuery data analytics to understand their website performance, but wants to move only some data into the cloud. Which environment should the organization use?

- A. Private cloud
- B. On-premises
- C. Multi-cloud
- D. Hybrid cloud

Answer: D

Explanation:

The assumption should be made that there is still a private network involved. Hybrid clouds always include a private cloud and are typically managed as one entity. Multi-clouds always include more than one public cloud service, which often perform different functions.

NEW QUESTION 86

- (Topic 3)

An organization provides a loyalty program for its customers. It recently partnered with other businesses so that customers can get loyalty points at a range of other stores.

Why should the organization use application programming interfaces (APIs)?

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

Answer: D

NEW QUESTION 87

- (Topic 3)

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

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

Answer: D

NEW QUESTION 91

- (Topic 3)

An organization wants to add a new function to their application. They want to write the code and let the public cloud provider handle the infrastructure. Which infrastructure solution should they use?

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

Answer: C

Explanation:

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

NEW QUESTION 92

- (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 93

- (Topic 3)

An organization operates their entire IT infrastructure from Google Cloud. What should they do to prepare for data breaches?

- A. Reduce reliance on multi-factor authentication
- B. Data security is Google's responsibility, so preparation is minimal
- C. Create an incident plan to mitigate impacts
- D. Strengthen their data center perimeter security

Answer: C

NEW QUESTION 95

- (Topic 3)

An online retail organization wants to optimize their service. What is an example of unstructured data that they can use to make decisions?

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

Answer: A

Explanation:

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

NEW QUESTION 100

- (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 101

- (Topic 3)

A cloud-native organization is not meeting their service level objective (SLO) but has not exhausted their error budget. What should the organization prioritize?

- A. Innovation to improve user experience
- B. Hardware reliability to improve availability
- C. Stability to avoid prolonged user downtime
- D. Speed to release new features

Answer: C

Explanation:

Both Devs and SRE team must ensure that the error budget does not become exhausted. To avoid it, releases have to stop for the time being until the error budget resets. The team would have to reprioritize to focus on reliability to get it back to an acceptable state.

NEW QUESTION 105

- (Topic 3)

An organization wants a cost-effective relational database. Which Google Cloud service should the organization use?

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

Answer: C

NEW QUESTION 110

- (Topic 3)

An organization has decided to modernize their applications in the cloud to keep up with their customers' needs. What may have prompted this business decision?

- A. Their on-premises applications only autoscale to meet demand.
- B. They want to change from a pay-as-you-go model to a capital expenditure model.
- C. Their source code changes erroneously without developer interaction.
- D. Their on-premises applications take months to update and deploy.

Answer: D

NEW QUESTION 111

- (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 114

- (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 115

- (Topic 3)

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

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

Answer: B

NEW QUESTION 118

- (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 121

- (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 126

- (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 131

- (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 132

- (Topic 3)

An e-commerce organization is reviewing their cloud data storage.

What type of raw data can they store in a relational database without any processing?

- A. Product inventory
- B. Product photographs
- C. Instructional videos
- D. Customer chat history

Answer: A

NEW QUESTION 136

- (Topic 3)

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

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

Answer: D

NEW QUESTION 138

- (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 139

- (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 143

- (Topic 3)

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

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

Answer: A

Explanation:

Virtual machines - you can install customized OS Containers - about applications

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

NEW QUESTION 145

- (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 150

- (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 152

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 156

- (Topic 3)

An organization wants its users to validate a series of new features for their app. Why should they use App Engine?

- A. Because their app is containerized and enabled by microservices
- B. Because the updated app will only include new features
- C. To run different versions of the app for different users
- D. To run different versions of the app for the same user

Answer: C

NEW QUESTION 159

- (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 161

- (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 165

- (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 167

- (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 171

- (Topic 3)

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

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

Answer: B

NEW QUESTION 173

- (Topic 3)

An organization wants to introduce a new image recognition login system. What should the organization do to follow SRE principles?

- A. Roll out the new system to a subset of employees to test it out
- B. Roll out the new system to all employees to collect as much data as possible
- C. Avoid rolling out the new system because it may have security flaws
- D. Avoid rolling out the new system because it may violate privacy policy

Answer: A

NEW QUESTION 178

- (Topic 3)

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 182

- (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 185

- (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 188

- (Topic 3)

What does Cloud Debugger help an organization do?

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

Answer: B

Explanation:

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

NEW QUESTION 190

- (Topic 2)

You are working with a user to set up an application in a new VPC behind a firewall and it is noticed that the user is concerned about data egress. Therefore, to provide assistance you want to configure the fewest open egress ports. Which of the following statements is correct?

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

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

Answer: B

Explanation:

Implied rules Every VPC network has two implied firewall rules. These rules exist, but are not shown in the Cloud Console:

Implied allow egress rule. An egress rule whose action is allow, destination is 0.0.0.0/0, and priority is the lowest possible (65535) lets any instance send traffic to any destination, except for traffic blocked by Google Cloud. A higher priority firewall rule may restrict outbound access. Internet access is allowed if no other firewall rules deny outbound traffic and if the instance has an external IP address or uses a Cloud NAT instance. For more information, see Internet access requirements.

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

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

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

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

NEW QUESTION 192

- (Topic 2)

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

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

Answer: BC

NEW QUESTION 195

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 197

- (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 198

- (Topic 2)

A large organization is struggling to manage their cloud costs effectively. They want to increase vis-ibility into cloud costs. Which cost management approach should the organization use?

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

Answer: A

Explanation:

Because cross-team partnerships are part of the visibility cost management strategy.

https://wa.aws.amazon.com/wat.question.COST_1.en.html

NEW QUESTION 202

- (Topic 2)

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

- About 12 days.
- About 2 years.
- About 100 days.
- About 24 hours.

Answer: C

Explanation:

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

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

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

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

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

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

NEW QUESTION 204

- (Topic 2)

Google Cloud Platform (GCP) provides three main compliance resource webpages. What are they? (Select Three Answer)

- Compliance Reports Manager
- Support Hub
- Compliance Offerings
- GDPR Home Page
- TechCentral

Answer: ACD

Explanation:

Compliance Reports Manager, GDPR Home Page, Compliance Offerings GCP provides three main compliance resource webpages

Compliance Reports Manager

– <https://cloud.google.com/security/compliance/compliance-reports-manager>

Compliance Reports Manager

Google Cloud's industry-leading security, third-party audits and certifications, documentation, and contract commitments help support your compliance. Compliance reports manager provides you with easy, on-demand access to these critical compliance resources, at no additional cost. Key resources include our latest ISO/IEC certificates, SOC reports, and self assessments.

Select resources may require sign-in with your Google Cloud or Google Workspace account. If you would like to access previous reports please reach out to [support](#) for more information. Anything marked "Google Confidential Information" is shared subject to the confidentiality obligations described in the customer or partner agreement(s) covering Cloud Services. Please contact your sales representative for permission to share confidential resources outside of your organization with customers or other third parties not expressly permitted by your agreement.

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

Compliance offerings

To help you with compliance and reporting, we share information, best practices, and easy access to documentation. Our products regularly undergo independent verification of security, privacy, and compliance controls, achieving certifications against global standards to earn your trust. We're constantly working to expand our coverage. This site contains information about Google's certifications and compliance standards it satisfies as well as general information about certain region or sector-specific regulations.

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

NEW QUESTION 208

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

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

Answer: B

Explanation:

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

NEW QUESTION 210

- (Topic 2)
 An application has become very popular and the number of requests/users is increasing quickly. There is a meeting to figure out how to scale the systems so that they can accept user requests and still have the capacity to spare. What is the preferred option?

- A. Circular Scaling takes a round-robin approach to allocate and destroy VMs.
- B. Triangular Scaling takes an automated average of Cost, Effort, and Time.
- C. Vertical Scaling
- D. Horizontal Scaling

Answer: D

Explanation:

Horizontal scaling, also called scaling out, adds new VMs to increase application capacity.

NEW QUESTION 215

- (Topic 2)
 A developer in your IT team is cheating 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 nam
- B. Look for similar bucket names that have a "." in it.
- C. Previously deleted bucket names in the same project cannot be reuse
- 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 name are case insensitive- look for bucket name 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 220

- (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 In-tervention.
- 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 224

- (Topic 2)

Which of the followings are core components of Anthos?

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

Answer: D

Explanation:

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

NEW QUESTION 226

- (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 230

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 234

- (Topic 2)

When you update the function in firebase by deploying updated code, instances for older versions are cleaned up along with build artifacts in and replaced by new instances.

- A. Google Cloud console.
- B. Storage and Container Registry.
- C. Container Registry repository.
- D. None of the Above

Answer: B

Explanation:

Container Registry is a single place for your team to manage Docker images, perform vulnerability analysis, and decide who can access what with fine-grained access control

NEW QUESTION 237

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 239

- (Topic 2)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 241

- (Topic 2)

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

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

Answer: C

Explanation:

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

Storage Transfer Service is a product that enables you to:

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

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

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

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

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

NEW QUESTION 244

- (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 248

- (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 250

- (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 253

- (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 257

- (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 258

- (Topic 2)

You have deployed a new public web application that allows users to register and login with email ids, phone numbers, or user ids. You are seeing some unusual activity with user registrations and logins from a few IPs. A large number of accounts were created very quickly. Logins are also happening quickly thereafter from these new accounts. Different parts of the application are being explored, all of which are putting a heavy load on the application. What could be a problem and how can you solve it?

- A. A hacker group has hired a bunch of people to create accounts and manually use the system
- B. Use Cloud Asset Inventory to see if there have been changes in the inventory
- C. Bots are creating accounts and then using the
- D. Use Google Cloud's Web App and API Protection (WAAP).
- E. Bots are creating accounts and then using the
- F. Use Identity-Aware Proxy to restrict the users to known users.
- G. Automated testing tools might still be running and creating account
- H. Use Identity-Aware Proxy to restrict the users to known users.

Answer: B

Explanation:

Bots attacking the application is the most likely scenario in this case. Using WAAP is the right protection plan: Anti-DDoS, anti-bot, WAF, and API protection help you protect against new and existing threats while helping you keep your apps and APIs compliant and continuously available.

<https://cloud.google.com/solutions/web-app-and-api-protection>

NEW QUESTION 259

- (Topic 2)

Which of the following is true while creating a boot persistent disk from a snapshot.

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

Answer: D

Explanation:

When you create a virtual machine (VM) instance, you must also create a boot disk for the VM. You can use a public image, a custom image, or a snapshot that was taken from another boot disk. When you create a boot disk, limit the disk size to 2 TB to account for the limitations of MBR partitioning.

Compute Engine automatically creates a boot persistent disk when you create an instance. If you require additional data storage space for your instances, add one or more secondary instance storage options.

You might need to create a standalone boot persistent disk and attach it to an instance later, or resize a boot persistent disk to improve performance and add more space for additional applications or operating system files. That process is described in Add or resize a persistent disk.

As a best practice, do not use regional persistent disks for boot disks. In a failover situation, they do not force-attach to a VM.

After you create a snapshot of a boot persistent disk, you can apply data from that snapshot to new persistent disks. It is only possible to apply data from a

snapshot when you first create a persistent disk. You cannot apply a snapshot to an existing persistent disk, or apply a snapshot to persistent disks that belong to a different project than that snapshot.

NEW QUESTION 263

- (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 266

- (Topic 2)

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

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

Answer: A

Explanation:

Folders for a related group of projects are the recommended approach.

-> A flat structure under the organization node is possible on Google Cloud, but it is not recommended. It becomes tougher to manage.

-> Projects cannot have sub-projects; there can only be resources within Projects.

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

NEW QUESTION 267

- (Topic 2)

You're negotiating SLAs with a customer. You have communicated that there will be a 99.99% (four 9s) availability for the service you are providing. Every aspect of the service is under your control. They want to modify the reliability to 99.999% (five 9s). What do you tell them? (Choose two answers)

- A. Yes, that could be possible
- B. If yes, there will be a significantly higher charge because the effort is significantly higher too.
- C. Yes, that is possible, but there will be an additional charge of 9% for the service because that is the additional effort required.
- D. Yes, that is possible
- E. There is hardly any difference to provide another 0.009% availability.
- F. Ask them for the reasonable downtime they are willing to absorb
- G. If it is more than 60 minutes in an entire year, explain how the current SLA meets that requirement.

Answer: AD

Explanation:

In many cases, customers might not know the implications of the 9s with respect to scheduled maintenance, upgrades, etc. It's possible that they are holding unnecessary expectations that significantly exceed their requirements.

-> Even though 0.0009% increase it looks like a small increment, an addition of a single 9 reduces the possible downtime by 10 times. So the effort is often much greater.

Reference link- https://en.wikipedia.org/wiki/High_availability

NEW QUESTION 269

- (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 274

- (Topic 2)

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

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

Answer: CDE

Explanation:

Sustained, Committed and Preemptible

vCPU and memory usage for each of these categories can receive discounts VM vCPU and memory usage for each of these categories can receive discounts Sustained-use discounts—Google offers up to 30% off for workloads that run for most of the billing month on GCP services.

Committed-use discounts—users can save up to 57% by committing to use an instance for a certain time period, with no upfront payment and with the flexibility to change instances during the commitment period.

Preemptible VMs—similar to the concept of AWS spot instances, Google offers up to 79% off for Virtual Machines that may be shut down at any time and replaced by others.

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

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

NEW QUESTION 279

- (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 testers 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 testers be permanently hired to plug the vulnerability issue.
- F. Create a Group with the permissions required to do the test and record their input
- 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 283

- (Topic 2)

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

- A. Recommend a move to serverless technologies which will scale automatically on demand.
- B. Add logging statements at multiple points in the application, build it, and deploy it
- C. Now new requests will give us information on latency in the logs.
- D. Check if the browsers used by the client are different from your
- E. If they are, that's most likely the issue
- F. Ensure that everybody uses the latest version of the browser that you are also using.
- G. Use Cloud Trace to collect latency data and track how requests propagate and why there is a delay.

Answer: D

Explanation:

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

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

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

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

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

NEW QUESTION 286

- (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 288

- (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 289

- (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 290

- (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 293

- (Topic 2)

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

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

Answer: D

Explanation:

Firebase Cloud Messaging:

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

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

Key capabilities of Firebase Cloud Messaging:

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

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

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

NEW QUESTION 295

- (Topic 2)

What cloud deployment model is generally deployed between organizations such as non- profits, hospitals or even enterprises that share similar requirements or interests?

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

Answer: B

Explanation:

Community Cloud – The cloud infrastructure is planned for selective use by a particular community of consumers from organizations that have mutual interests like security needs, policy, and compliance considerations.

Reference link- https://csrc.nist.gov/glossary/term/community_cloud

NEW QUESTION 298

- (Topic 2)

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 303

- (Topic 2)

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

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

Answer: BC

Explanation:

Customer-managed encryption keys for Cloud BigTable.

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

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

Features

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

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

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

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

NEW QUESTION 304

- (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 in-formation 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 308

- (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 313

- (Topic 2)

When creating machine learning models, a key initial step is to identify the type of model required. One of these is the classification model. Which of these statements define a classification model?

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

Answer: A

Explanation:

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

NEW QUESTION 315

- (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 320

- (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 321

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 326

- (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 331

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 332

- (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 337

- (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 342

- (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 347

- (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 351

- (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 354

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 359

- (Topic 1)

A fitness band company is continuously ingesting data from millions of its consumers. Different kinds of data based on time, like location, heartbeat rate, temperature, movement, etc. are connect-ed. They need a high throughput database that can write data very fast. Since their users are spread across the world, they need the database to be geographically scalable. Consumers also want to see near-real-time visualizations of their activities. Which of these databases would be a good fit?

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

Answer: B

Explanation:

Bigtable is the best suited for time series data. It also has high read-write throughput and ability to scale globally.

NEW QUESTION 360

- (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 362

- (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 365

- (Topic 1)

An organization has created an ecommerce website. What data on this website would be considered structured data?

- A. Product photographs
- B. Product reviews
- C. Product descriptions
- D. Product ratings score

Answer: D

Explanation:

Because product ratings are structured because they are numerical scores.

NEW QUESTION 367

- (Topic 1)

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

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

Answer: B

Explanation:

It would be most straightforward to use AutoML Image. Put the images in Cloud Storage, point to it from AutoML, and start the model building process.

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

NEW QUESTION 372

- (Topic 1)

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

- A. Cloud Functions

- B. App Engine
- C. Compute Engine
- D. Anthos

Answer: D

Explanation:

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

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

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

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

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

NEW QUESTION 376

- (Topic 1)

Your team is using BigQuery as your central data warehouse. You are running a certain workload that you've run frequently over the last few days. It is a short, high capacity analytics workload. Which of the following would be an appropriate pricing model to use?

- A. There is no need for any pricing model the first 1 TB of query data processed per month is free.
- B. On-demand pricing
- C. Flex Slots
- D. Flat-rate reservations

Answer: C

Explanation:

Option A is Correct- BigQuery Flex Slots for cyclical workloads that require extra capacity, or for workloads that need to process a lot of data in a short time, and so would be less expensive to run using reserved slots for a short time.

NEW QUESTION 381

- (Topic 1)

What are the network requirements for Private Google Access?

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

Answer: B

Explanation:

Network requirements for Private Google Access:

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

NEW QUESTION 385

- (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 387

- (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 388

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

Cloud-Digital-Leader Practice Exam Features:

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

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