



Google

Exam Questions Cloud-Digital-Leader

Google Cloud Digital Leader exam

NEW QUESTION 1

- (Exam Topic 2)

You have experimented with Google Cloud using your own credit card and expensed the costs to your company. Your company wants to streamline the billing process and charge the costs of your projects to their monthly invoice. What should you do?

- A. Grant the financial team the IAM role of €Billing Account User€ on the billing account linked to your credit card.
- B. Change the billing account of your projects to the billing account of your company.
- C. Create a ticket with Google Billing Support to ask them to send the invoice to your company.
- D. Set up BigQuery billing export and grant your financial department IAM access to query the data.

Answer: B

Explanation:

To change the Cloud Billing account for a project, you need to be able to move a project from one Cloud Billing account to another. To accomplish this task, you need permissions adequate to unlink the project from the existing Cloud Billing account AND to link the project to the target Cloud Billing account. Roles with adequate permissions to perform this task: Project Owner or Project Billing Manager on the project, AND Billing Account Administrator or Billing Account User for the target Cloud Billing account

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

A Cloud Billing account is used to define who pays for a given set of resources, and it can be linked to one or more projects. Project usage is charged to the linked Cloud Billing account.

If you are a billing administrator on only one Cloud Billing account, new projects you create are automatically linked to your existing Cloud Billing account. If you create or have access to multiple Cloud Billing accounts, you can change the Cloud Billing account a project is billed to. This article describes how to change the Cloud Billing account for your project, as well as how to enable and disable billing for a project.

Reference link- <https://cloud.google.com/billing/docs/how-to/modify>

NEW QUESTION 2

- (Exam Topic 2)

Your customer is moving to Google Cloud. They have many teams, each working on many projects. How should they organize resources?

- A. Let each team have one shared Folder with multiple Projects within it so that there is a separation of concerns.
- B. Let each Project have one Folder so that there is a clear separation of concerns.
- C. Let each team have an Organization so that they can entirely manage themselves with their own identity.
- D. Let each team have one shared Project so that it is easy to manage.

Answer: A

Explanation:

The recommended approach is to have folders corresponding to teams/departments and they manage the projects within that.

-> Sharing a single project will cause a conflict of resources, billing, concerns, etc.

-> One folder per project is unnecessary overuse of abstraction/grouping.

-> Teams and projects in a company should ideally be centrally managed in a single Organization.

NEW QUESTION 3

- (Exam Topic 2)

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

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

Answer: D

Explanation:

Flavors of Apigee

Apigee comes in the following flavors:

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

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

NEW QUESTION 4

- (Exam Topic 2)

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

- A. Multi-Regional Storage

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

Answer: B

Explanation:

Coldline storage is a very low cost highly durable storage service for data archiving, online backup, and disaster recovery. Coldline storage is the best choice for data that you plan to access at most once a year due to its slightly lower availability, 90 day minimum storage duration cost for data access, and higher per operation costs.

Nearline and Coldline are for backup and archival storage and having the highest availability for both with 99.9 percent.

NEW QUESTION 5

- (Exam 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 6

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 7

- (Exam Topic 2)

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

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

Answer: AB

Explanation:

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

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

NEW QUESTION 8

- (Exam 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. Text Description automatically generated

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 9

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

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

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

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

You can manually import boot disk images to Compute

Engine, but one disk at a time.

Graphical user interface, text, application, email Description automatically generated <https://cloud.google.com/compute/docs/images>

NEW QUESTION 10

- (Exam 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 10

- (Exam 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.

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

Introduction to the Organization Policy

Send feedback

Service

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 15

- (Exam Topic 1)

Your organization needs to establish private network connectivity between its on-premises network and its workloads running in Google Cloud. You need to be able to set up the connection as soon as possible.

Which Google Cloud product or feature should you use?

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

Answer: A

Explanation:

Private Google Access for on-premises hosts provides a way for on-premises systems to connect to Google APIs and services by routing traffic through a Cloud VPN tunnel.

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

NEW QUESTION 16

- (Exam Topic 1)

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

What should you do?

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

Answer: C

Explanation:

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

NEW QUESTION 18

- (Exam 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 21

- (Exam Topic 1)

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 22

- (Exam Topic 1)

An organization currently stores its data on-premises and they receive different levels of traffic on their website every month. How could the organization benefit from modernizing their infrastructure with cloud technology?

- A. They can rely on the cloud provider for all website source code.
- B. Agile storage scalability.
- C. 100% service availability.
- D. They can shift from heavy operational expenditure to a capital expenditure model.

Answer: B

Explanation:

Organizations can scale in the cloud by paying for what they use, when they use it.

NEW QUESTION 27

- (Exam Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 29

- (Exam Topic 1)

Your organization uses Active Directory to authenticate users. Users' Google account access must be removed when their Active Directory account is terminated. How should your organization meet this requirement?

- A. Configure two-factor authentication in the Google domain
- B. Remove the Google account from all IAM policies
- C. Configure BeyondCorp and Identity-Aware Proxy in the Google domain
- D. Configure single sign-on in the Google domain

Answer: D

Explanation:

Configure single sign-on in the Google domain

Single sign-on: Whenever a user needs to authenticate, Google Cloud delegates the authentication to Active Directory by using the Security Assertion Markup Language (SAML) protocol. This delegation ensures that only Active Directory manages user credentials and that any applicable policies or multi-factor authentication (MFA) mechanisms are being enforced. For a sign-on to succeed.

Federating Google Cloud with Active Directory 🔖

[Send feedback](#)

This article describes how you can configure Cloud Identity or Google Workspace to use [Active Directory as IdP and authoritative source](#).

The article compares the logical structure of Active Directory with the structure used by Cloud Identity and Google Workspace and describes how you can map Active Directory forests, domains, users, and groups. The article also provides a [flowchart](#) that helps you determine the best mapping approach for your scenario.

This article assumes that you're familiar with Active Directory.

Implementing federation

Google Cloud uses [Google Identities](#) for authentication and access management. Manually maintaining Google identities for each employee can add unnecessary management overhead when all employees already have an account in Active Directory. By federating user identities between Google Cloud and your existing identity management system, you can automate the maintenance of Google identities and tie their lifecycle to existing users in Active Directory.

<https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-introduction>
 Reference Link- <https://cloud.google.com/architecture/identity/single-sign-on>

NEW QUESTION 31

- (Exam 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 32

- (Exam 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 36

- (Exam 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 41

- (Exam Topic 1)

Your company security team manages access control to production systems using an LDAP directory group. How is this access control managed in the Google Cloud production project?

- A. Assign the proper role to the Service Account in the project's IAM Policy
- B. Grant each user the roles/iam.serviceAccountUser role on a service account that exists in the Google Group.
- C. Assign the proper role to the Google Group in the project's IAM Policy.
- D. Create the project in a folder with the same name as the LDAP directory group.

Answer: C

Explanation:

Reference:

<https://cloud.google.com/blog/products/identity-security/achieving-identity-and-access-governance-on-google-c>
Text, letter Description automatically generated

NEW QUESTION 43

- (Exam 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 45

- (Exam Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 49

- (Exam 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 50

- (Exam Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 55

- (Exam Topic 1)

What conditions be true if a VM interface wants to send packets to the external IP addresses of Google APIs and services using Private Google Access?

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

Answer: A

Explanation:

A VM interface can send packets to the external IP addresses Google Access if all these conditions are met:

of Google APIs and services using Private

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

NEW QUESTION 59

- (Exam Topic 1)

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

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

Answer: AB

Explanation:

Graphical user interface, text Description automatically generated with medium confidence



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

NEW QUESTION 64

- (Exam 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.

Text Description automatically generated

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_u

NEW QUESTION 68

- (Exam 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:

Text Description automatically generated

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.

Text Description automatically generated

```
# Template for new folder & new project

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

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

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

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

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

NEW QUESTION 71

- (Exam 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 73

- (Exam 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 76

- (Exam 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 80

- (Exam Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 85

- (Exam 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 90

- (Exam Topic 1)

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

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

Answer: D

Explanation:

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

Text Description automatically generated

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

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

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

NEW QUESTION 93

- (Exam Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 94

- (Exam 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 `threerestricted.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 the `restricted.googleapis.com`

create multiple routes.

domain names, you only need one route (per domain). Otherwise, you'll need to

- 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 99

- (Exam Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 103

- (Exam Topic 1)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 107

- (Exam 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 108

- (Exam Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 113

- (Exam Topic 1)

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 118

- (Exam 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 120

- (Exam 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 122

- (Exam Topic 1)

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

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

Answer: D

Explanation:

Premium Support will have a 15-minute response time with 24/7 response for high & critical-impact issues. Graphical user interface, text, application, email Description automatically generated

Premium Support overview

[Send feedback](#)

This page explains the features of Premium Support.

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

This overview covers the following aspects of Premium Support:

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

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

NEW QUESTION 125

- (Exam 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 128

- (Exam Topic 1)

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

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

Answer: C

Explanation:

The pricing calculator can be used to give clear estimates of resource usage.

-> Running test loads is as closely indicative as using the pricing calculator.

-> There are no cloud cost references published, either by Google or CNCF. Even if some companies have published such info. It might not apply to you.

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

NEW QUESTION 129

- (Exam Topic 1)

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

What should you do?

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

Answer: D

Explanation:

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

NEW QUESTION 131

- (Exam Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 135

- (Exam 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 140

- (Exam Topic 1)

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

- A. Customers expect faster time to market for games.
- B. Employees expect source code changes to be deployed faster.
- C. Customers expect a personalized experience.

D. Employees expect more predictable data management spending.

Answer: C

Explanation:

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

NEW QUESTION 142

- (Exam 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 143

- (Exam Topic 1)

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

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

Answer: B

Explanation:

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

Reference:

<https://www.okta.com/blog/2021/07/what-is-an-identity-platform/#:-:text=An%20identity%20platform%20is%20Text%20Description%20automatically%20generated>

Text Description automatically generated

NEW QUESTION 144

- (Exam Topic 1)

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 146

- (Exam Topic 1)

Your organization needs to plan its cloud infrastructure expenditures. Which should your organization do?

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

Answer: A

Explanation:

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

NEW QUESTION 150

- (Exam Topic 1)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 151

- (Exam Topic 1)

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

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 152

- (Exam Topic 1)

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

- A. Choose the Standard network service tier.
- B. Choose the Premium network service tier.
- C. Deploy Cloud VPN.
- D. Deploy Cloud NAT.

Answer: A

Explanation:

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

Text, letter Description automatically generated

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

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

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

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

NEW QUESTION 155

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

NEW QUESTION 160

- (Exam 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:

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 162

- (Exam Topic 1)

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

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

Answer: B

Explanation:

Store the information in Secret Manager is a secure and convenient storage system for API keys, passwords, certificates, and other sensitive data. Secret Manager provides a central place and single source of truth to manage access, and audit secrets across Google Cloud.
<https://cloud.google.com/secret-manager>

NEW QUESTION 165

- (Exam 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 166

- (Exam 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 170

- (Exam 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 172

- (Exam Topic 3)

An organization delivers a proactive healthcare service. They want to efficiently and automatically collect patient data. What should the organization encourage the patients to do?

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

Answer: B

NEW QUESTION 176

- (Exam 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 177

- (Exam Topic 3)

How does a large hotel chain benefit from storing their customer reservation data in the cloud?

- A. On-premises hardware access to transaction data
- B. Real-time data transformation at scale within an on-premises database
- C. Real-time business transaction accuracy at scale
- D. Physical hardware access during peak demand

Answer: C

NEW QUESTION 182

- (Exam 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 183

- (Exam 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 187

- (Exam 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 191

- (Exam 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.

Text Description automatically generated

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-c>

NEW QUESTION 192

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 193

- (Exam 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?

- SLA - Service Level Agreement
- SLC - Service Level Contract
- SLO - Service Level Objective
- 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 194

- (Exam Topic 2)

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

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

Answer: A

Explanation:

Because Dataproc is used to run Hadoop/Spark workloads

NEW QUESTION 196

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 200

- (Exam 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 204

- (Exam 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 207

- (Exam 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 210

- (Exam 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 statement is correct?

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

Answer: B

Explanation:

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

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

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

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

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

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

NEW QUESTION 215

- (Exam 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 216

- (Exam 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 hap-pening quickly thereafter from these new accounts. Different parts of the application are being ex-plored, 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 syste
- 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 re-strict 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 217

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 220

- (Exam Topic 2)

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

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

Answer: C

Explanation:

Google provides a page that brings together everything needed around support. Its called the Support Hub

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

NEW QUESTION 225

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

Reference link

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

NEW QUESTION 227

- (Exam 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 231

- (Exam Topic 2)

What are the different storage & database services in GCP? Which is Google cloud storage and database below the option

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

Answer: E

Explanation:

Graphical user interface, text, application Description automatically generated

Google Cloud offers 9 storage and database options namely:

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

NEW QUESTION 232

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

Instance groups

Send feedback 

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

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

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

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

NEW QUESTION 235

- (Exam 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.

- This statement is undefined.
- The above statement is partially true.
- The above statement is completely false.
- The above statement is completely true.

Answer: D

Explanation:

Firestore Cloud Messaging:

Firestore 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 Firestore 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 239

- (Exam Topic 2)

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

- Recommend a move to serverless technologies which will scale automatically on demand.
- Add logging statements at multiple points in the application, build it, and deploy i
- Now new requests will give us information on latency in the logs.
- Check if the browsers used by the client are different from your
- If they are, that's most likely the issue.Ensure that everybody uses the latest version of the browser that you are also using.
- 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 server-ful technologies.

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

NEW QUESTION 243

- (Exam Topic 2)

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

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

Answer: D

Explanation:

Preemptible instances function like normal instances but have the following limitations:

- > Compute Engine might stop preemptible instances at any time due to system events. The probability that Compute Engine will stop a preemptible instance for a system event is generally low, but might vary from day to day and from zone to zone depending on current conditions.
- > Compute Engine always stops preemptible instances after they run for 24 hours. Certain actions reset this 24-hour counter.
- > Preemptible instances are finite Compute Engine resources, so they might not always be available.
- > Preemptible instances can't live migrate to a regular VM instance, or be set to automatically restart when there is a maintenance event.
- > Due to the above limitations, preemptible instances are not covered by any Service Level Agreement (and, for clarity, are excluded from the Compute Engine SLA).
- > The Google Cloud Free Tier credits for Compute Engine do not apply to preemptible instances.

NEW QUESTION 246

- (Exam Topic 2)

All Google Cloud Platform services are associated with a project that is used to provide what functions?

- A. Manage Container Deployments
- B. Enable Services and APIs
- C. Manage DNS Services
- D. None of the Above

Answer: B

Explanation:

Enable Services and APIs

Graphical user interface, text Description automatically generated

A project organizes all your Google Cloud resources. A project consists of a set of users; a set of APIs; and billing, authentication, and monitoring settings for those APIs. So, for example, all of your Cloud Storage buckets and objects, along with user permissions for accessing them, reside in a project. You can have one project, or you can create multiple projects and use them to organize your Google Cloud resources, including your Cloud Storage data, into logical groups.

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

NEW QUESTION 251

- (Exam 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.

Graphical user interface, text, application Description automatically generated

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 256

- (Exam Topic 2)

In discussions with a prospective customer who wants to move to Google Cloud to make use of the latest, scalable technologies available therein, you learn that there are very strict regulations concerning the storage of data. They only have the approval to store it in their current private data center. What would you advise them?

- A. Retain on-premise itself those portions of data and compute which are under regulation
- B. Take advantage of all the other cloud capabilities for remaining work-loads.
- C. It is too risky to touch anything in such a scenario
- D. It is best to remain entirely on-premise.

- E. Regulations are guideline
- F. As long as the data remains encrypted, you can move it anywhere.
- G. Petition the government for changes to such regulations as all industries are moving to the public cloud. Then, when the regulations are eased, move to Google Cloud.

Answer: A

Explanation:

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

NEW QUESTION 260

- (Exam 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 262

- (Exam Topic 2)

Which of the following methods should you use when you require a dynamic way of provisioning VMs on Compute Engine when it is observed that the exact specifications will be in a dedicated configuration file and you want to follow Google's recommended practices.

- A. Managed Instance Group
- B. Deployment Manager
- C. Cloud Composer
- D. Unmanaged Instance Group

Answer: B

Explanation:

The question is about a dynamic way to provision VM, it can be achieved by a Deployment manager or by using terraform. MIG is creating multiple machines based on templates by load balancing

NEW QUESTION 267

- (Exam Topic 2)

If you increase the size of a subnet in a custom VPC network, the IP addresses of virtual machines already on that subnet might be affected. Which options are Correct.

- A. False
- B. None of the above
- C. True
- D. Not Defined by Google Cloud Platform

Answer: A

Explanation:

You can dynamically increase the size of a subnet in a custom network by expanding the range of IP addresses allocated to it. Doing that doesn't affect already configured VMs.

NEW QUESTION 272

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 275

- (Exam 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 280

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

Cloud SQL Key Terms:

Cloud SQL instance

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

Database instance

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

High availability

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

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

Failover

A failover is when Cloud SQL switches serving from the original primary instance to the standby instance. Autofailover is a mechanism that automatically triggers failover when a Cloud SQL instance didn't issue a heartbeat in the previous interval.

Standby instances

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

Clone

When you clone a Cloud SQL instance, you create a new instance that is a copy of the source instance, but is completely independent. After cloning is complete, changes to the source instance are not reflected in the clone, and changes in the clone are not reflected in the source instance.

Replication

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

Read replica

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

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

Source server

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

Cloud SQL Auth proxy client

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

NEW QUESTION 281

- (Exam Topic 2)

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

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

Answer: ACD

Explanation:

Compliance Reports Manager, GDPR Home Page, Compliance Offerings

GCP provides three main compliance resource webpages

Compliance Reports Manager – <https://cloud.google.com/security/compliance/compliance-reports-manager> Text, timeline Description automatically generated

Compliance 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.

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

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.

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 286

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

Table Description automatically generated with medium confidence

Standard storage

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

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

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

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

Availability

The availability of Standard storage data is:

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

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

NEW QUESTION 288

- (Exam Topic 2)

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

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

Answer: D

Explanation:

Compute Engine provides graphics processing units (GPUs) that you can add to your virtual machines (VMs). You can use these GPUs to accelerate specific workloads on your VMs such as machine learning and data processing.

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

NEW QUESTION 290

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 294

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

Reference link

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

NEW QUESTION 298

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 301

- (Exam Topic 2)

Cloud SQL is a fully-managed relational database service for MySQL, PostgreSQL and SQL servers, keeping Cloud SQL Google Cloud Service in mind, which of the following statements is/are correct?

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

Answer: D

Explanation:

Cloud SQL

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

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

Fully managed

Cloud SQL automatically ensures your databases are reliable, secure, and scalable so that your business

continues to run without disruption. Cloud SQL automates all your backups, replication, encryption patches, and capacity increases—while ensuring greater than 99.95% availability, anywhere in the world.

Integrated

Access Cloud SQL instances from just about any application. Easily connect from App Engine, Compute Engine, Google Kubernetes Engine, and your workstation. Open up analytics possibilities by using BigQuery to directly query your Cloud SQL databases.

Reliable

Easily configure replication and backups to protect your data. Go further by enabling automatic failover to make your database highly available. Your data is automatically encrypted, and Cloud SQL is SSAE 16, ISO 27001, and PCI DSS compliant and supports HIPAA compliance.

Easy migrations to Cloud SQL

Database Migration Service (DMS) makes it easy to migrate your production databases to Cloud SQL with minimal downtime. This serverless offering eliminates the manual hassle of provisioning, managing, and monitoring migration-specific resources. DMS leverages the native replication capabilities of MySQL and PostgreSQL to maximize the fidelity and reliability of your migration. And it's available additional charge for native like-to-like migrations to Cloud SQL.

NEW QUESTION 306

- (Exam Topic 2)

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

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

Answer: C

Explanation:

In-App Messaging

Engage active app users with contextual messages.

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

NEW QUESTION 308

- (Exam 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 309

- (Exam 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 313

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 317

- (Exam 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 322

- (Exam Topic 3)

How would a global organization benefit from managing their data with Cloud Spanner?

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

Answer: B

Explanation:

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

NEW QUESTION 327

- (Exam 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 329

- (Exam 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 332

- (Exam Topic 3)

An organization needs to categorize a large group of photographs using pre-trained machine learning. Which Google Cloud product or service should the organization use?

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

Answer: A

Explanation:

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

NEW QUESTION 334

- (Exam 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

NEW QUESTION 335

- (Exam 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 338

- (Exam Topic 3)

How does Cloud SQL help organizations create business insights?

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

Answer: A

Explanation:

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

NEW QUESTION 339

- (Exam Topic 3)

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

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

Answer: D

Explanation:

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

<https://cloud.google.com/vertex-ai>

NEW QUESTION 341

- (Exam 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 345

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 349

- (Exam Topic 3)

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

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

Answer: B

NEW QUESTION 351

- (Exam Topic 3)

Why is data stored in Google Cloud secure and private?

- A. Data is encrypted by the Security Command Center
- B. Data is encrypted by Cloud Data Loss Prevention

- C. Data is encrypted by default
- D. Data is encrypted when an appropriate tag is applied

Answer: C

Explanation:

<https://cloud.google.com/docs/security/encryption/default-encryption#:~:text=Google%20uses%20the%20Adva>

NEW QUESTION 356

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 360

- (Exam 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 361

- (Exam 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 365

- (Exam 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 368

.....

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](#)