



Google

Exam Questions Cloud-Digital-Leader

Google Cloud Digital Leader exam

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

- (Exam Topic 2)

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

- A. Since the app is anyways going to be public, create permanent credentials for the UX testers that they can conveniently use each time.
- B. It would be a security issue to have users come and go
- C. Recommend that the testers be permanently hired to plug the vulnerability issue.
- D. It would be a security issue to have users come and go
- E. Recommend that the testers be permanently hired to plug the vulnerability issue.
- F. Create a Group with the permissions required to do the test and record their input
- G. When users arrive each week, add them to the group and after the testing period, remove them from the group.

Answer: D

Explanation:

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

NEW QUESTION 2

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 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)

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

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

Answer: C

Explanation:

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

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

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

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

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

Reference link

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

NEW QUESTION 5

- (Exam Topic 2)

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

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

Answer: CDE

Explanation:

Instance is alive for no more than 24 hours, Can be pre-empted with a 30 second notice, Discounted Significantly.

Preemptible VM is an instance that you can create and run at a lower cost than normal instances.

However, Compute Engine might stop (pre-empt) these instances if it requires access to those resources for other tasks. Preemptible instances are excess Compute Engine capacity, so their availability varies with usage.

Live at most 24 hours Can be pre-empted with a 30 second notification via API and are Discounted significantly

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

NEW QUESTION 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)

A customer has contacted you about migrating to Google Cloud. The customer would like to migrate their data from on premises as soon as possible. They don't have the budget to rewrite code, and they want the most direct route. What migration option should suggest to the customer?

- A. None, since the customer is not cloud native ready.
- B. Rip and Replace
- C. Lift and Shift
- D. Improve and Move

Answer: C

Explanation:

With Lift and Shift migrations, the customer could move workloads from a source environment to a target environment with few or no modifications or refactoring

Lift and shift

In a lift and shift migration, you move workloads from a source environment to a target environment with minor or no modifications or refactoring. The modifications you apply to the workloads to migrate are only the minimum changes you need to make in order for the workloads to operate in the target environment.

A lift and shift migration is ideal when a workload can operate as-is in the target environment, or when there is little or no business need for change. This migration is the type that requires the least amount of time because the amount of refactoring is kept to a minimum.

There might be technical issues that force a lift and shift migration. If you cannot refactor a workload to migrate and cannot decommission the workload, you must use a lift and shift migration. For example, it can be difficult or impossible to modify the source code of the workload, or the build process isn't straightforward so producing new artifacts after refactoring the source code might not be possible.

Lift and shift migrations are the easiest to perform because your team can continue to use the same set of tools and skills that they were using before. These migrations also support off-the-shelf software. Because you migrate existing workloads with minimal refactoring, lift and shift migrations tend to be the quickest, compared to improve and move or remove and replace migrations.

On the other hand, the results of a lift and shift migration are non-cloud-native workloads running in the target environment. These workloads don't take full advantage of cloud platform features, such as horizontal scalability, fine-grained pricing, and highly managed services.

<https://cloud.google.com/architecture/migration-to-gcp-getting-started>

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 2)

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

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

Answer: BC

NEW QUESTION 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



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 10

- (Exam Topic 1)

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

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

Answer: D

Explanation:

Bare Metal Solution for Oracle

Bring your Oracle workloads to Google Cloud with Bare Metal Solution and jumpstart your cloud journey with minimal risk.

- Continue to run any version, any feature set, any database option, and any customizations (patchsets)
- Enterprise-grade deployment platform
- High availability with Oracle RAC
- Works with any application, any Oracle versions
- All your existing investment in tooling and best practices will work as is

NEW QUESTION 11

- (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 16

- (Exam Topic 1)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 19

- (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 22

- (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 25

- (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 26

- (Exam Topic 1)

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

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

Answer: C

Explanation:

you can get a 57% discount by agreeing to commit to a 3-year contract. Any usage over the commitment will just be billed at the standard rate. Since they only

need 300 CPUs 30% of the time, will pick answer C so that we are not paying usage off 300 CPUs all of the time. This gives us a discount of 57% for 200 CPU's, huge cost savings.

NEW QUESTION 28

- (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 30

- (Exam Topic 1)

Your organization needs to allow a production job to have access to a BigQuery dataset. The production job is running on a Compute Engine instance that is part of an instance group. What should be included in the IAM Policy on the BigQuery dataset?

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

Answer: C

Explanation:

When an identity calls a Google Cloud API, BigQuery requires that the identity has the appropriate permissions to use the resource. You can grant permissions by granting roles to a user, a group, or a service account.
 Reference link- <https://cloud.google.com/bigquery/docs/access-control>

NEW QUESTION 31

- (Exam Topic 1)

Your organization consists of many teams. Each team has many Google Cloud projects. Your organization wants to simplify the management of identity and access policies for these projects. How can you group these projects to meet this goal?

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

Answer: C

Explanation:

Text Description automatically generated with medium confidence

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

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

NEW QUESTION 34

- (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 35

- (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 40

- (Exam Topic 1)

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

Which Google Cloud product should your organization choose?

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

Answer: D

Explanation:

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

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

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

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

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

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

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

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

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

NEW QUESTION 43

- (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 48

- (Exam Topic 1)

Your team is working on building a machine learning model. There are a bunch of terminologies that are being used. What is an "instance" or an "example"?

- A. An input variable is used in making prediction
- B. E.
- C. number of rooms in a house price prediction model.
- D. One row of a dataset containing one or more input columns and possibly a prediction result.
- E. An answer for a prediction task, either the answer produced by a machine learning system or the right answer supplied in training data
- F. E.
- G. image contains a "cat".
- H. The "knobs" that you tweak during successive runs of training a model
- I. E.
- J. learning rate

Answer: B

Explanation:

One row of a dataset containing one or more input columns and possibly a prediction result. A picture containing timeline Description automatically generated

- **Instance:** The thing about which you want to make a prediction. For example, the instance might be a web page that you want to classify as either "about cats" or "not about cats".
- **Label:** An answer for a prediction task either the answer produced by a machine learning system, or the right answer supplied in training data. For example, the label for a web page might be "about cats".
- **Feature:** A property of an instance used in a prediction task. For example, a web page might have a feature "contains the word 'cat'".
- **Feature Column:** A set of related features, such as the set of all possible countries in which users might live. An example may have one or more features present in a feature column. "Feature column" is Google-specific terminology. A feature column is referred to as a "namespace" in the VW system (at Yahoo/Microsoft), or a **field**.
- **Example:** An instance (with its features) and a label.
- **Model:** A statistical representation of a prediction task. You train a model on examples then use the model to make predictions.

<https://developers.google.com/machine-learning/guides/rules-of-ml#terminology>

NEW QUESTION 52

- (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 56

- (Exam Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 59

- (Exam Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 60

- (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 65

- (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 67

- (Exam Topic 1)

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

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

Answer: C

Explanation:

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

References:

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

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

NEW QUESTION 71

- (Exam Topic 1)

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

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

Answer: B

Explanation:

Spanner, Cloud SQL, BigQuery

Spanner- Cloud Spanner is a fully managed, mission-critical, relational database service that offers transactional consistency at global scale, automatic, synchronous replication for high availability, and support for two SQL Google Standard SQL and PostgreSQL.

Cloud SQL- Cloud SQL is a fully-managed database service that helps you set up, maintain, manage, and administer your relational databases on Google Cloud Platform.

BigQuery- Google BigQuery is a cloud-based Architecture and provides exceptional performance as it can auto-scale up and down based on the data load and performs data analysis efficiently. On the other hand, SQL Server is based on client-server architecture and has fixed performance throughout unless the user scales it manually.

NEW QUESTION 74

- (Exam Topic 1)

Your ed-tech start-up was originally launched in a small geography. Any user sign-ups, course progress, tests taken, etc. are captured on a self-managed MySQL

database. Every user generates many such transactions. Now you're taking the application globally and preparing for a much larger influx of users from all over the world. The existing MySQL server is unlikely to be able to scale. Which convenient option can be considered?

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

Answer: B

Explanation:

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

NEW QUESTION 79

- (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 80

- (Exam Topic 1)

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

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

Answer: C

Explanation:

Reference:

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

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

NEW QUESTION 83

- (Exam Topic 1)

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

What should your organization do?

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

Answer: C

Explanation:

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

Direct Peering overview

[Send feedback](#)

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

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

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

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

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

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

NEW QUESTION 87

- (Exam Topic 1)

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 88

- (Exam Topic 1)

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

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

Answer: B

Explanation:

Migrate for Compute Engine enables you to lift and shift workloads at scale to Google Cloud Compute Engine with minimal changes and risk.

Reference: <https://dataintegration.info/simplify-vm-migrations-with-migrate-for-compute-engine-as-a-service>

NEW QUESTION 89

- (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%](https://www.okta.com/blog/2021/07/what-is-an-identity-platform/#:~:text=An%20identity%20platform%20is%20)

Text Description automatically generated

NEW QUESTION 93

- (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 97

- (Exam Topic 1)

Your organization is developing a mobile app and wants to select a fully featured cloud-based compute platform for it. Which Google Cloud product or feature should your organization use?

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

Answer: B

Explanation:

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

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

NEW QUESTION 101

- (Exam Topic 1)

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

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

Answer: AD

Explanation:

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

NEW QUESTION 105

- (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 108

- (Exam Topic 1)

What is the difference between Standard and Coldline storage?

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

Answer: D

Explanation:

Reference: <https://www.msp360.com/resources/blog/google-cloud-nearline-storage-vs-coldline-vs-standard/> Google Cloud Coldline is a new cold-tier storage for archival data with access frequency of less than once per year. Unlike other cold storage options, Nearline has no delays prior to data access, so now it is the leading solution among competitors.

NEW QUESTION 109

- (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 113

- (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 115

- (Exam Topic 1)

The operating systems of some of your organization's virtual machines may have a security vulnerability. How can your organization most effectively identify all virtual machines that do not have the latest security update?

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

Answer: A

Explanation:

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

Reference link:

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

NEW QUESTION 117

- (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 119

- (Exam Topic 3)

An organization wants to use Apigee to manage all their application programming interfaces (APIs). What will Apigee enable the organization to do?

- A. Increase application privacy
- B. Measure and track API performance Most Voted
- C. Analyze application development speed
- D. Market and sell APIs

Answer: B

Explanation:

Apigee's API Monitoring enables you to track your APIs to make sure they are up and running correctly. API Monitoring provides near real-time insights into API traffic and performance, to help you quickly diagnose and solve issues as they arise.

Apigee works with APIs not necessarily applications. It allows organizations to gain actionable insights across the entire API value chain and monetize API products and maximize the business value of digital assets. <https://cloud.google.com/apigee#section-11>

NEW QUESTION 124

- (Exam Topic 3)

An online retail organization wants to optimize their service.

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

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

Answer: A

Explanation:

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

NEW QUESTION 126

- (Exam Topic 3)

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

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

Answer: B

NEW QUESTION 129

- (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 131

- (Exam Topic 2)

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

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

Answer: D

Explanation:

Kubernetes vs. Docker

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

Docker lets you put everything you need to run your application into a box that can be stored and opened when and where it is required. Once you start boxing up your applications, you need a way to manage them; and that's what Kubernetes does.

Kubernetes is a Greek word meaning 'captain' in English. Like the captain is responsible for the safe journey of the ship in the seas, Kubernetes is responsible for carrying and delivering those boxes safely to locations where they can be used.

- Kubernetes can be used with or without Docker.
- Docker is not an alternative to Kubernetes, so it's less of a "Kubernetes vs. Docker" question. It's about using Kubernetes with Docker to containerize your applications and run them at scale.
- The difference between Docker and Kubernetes relates to the role each play in containerizing and running your applications.
- Docker is an open industry standard for packaging and distributing applications in containers.
- Kubernetes uses Docker to deploy, manage, and scale containerized applications.

NEW QUESTION 134

- (Exam Topic 2)

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

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

Answer: CDE

Explanation:

Sustained, Committed and Preemptible

vCPU and memory usage for each of these categories can receive discounts VM vCPU and memory usage for each of these categories can receive discounts

Sustained-use discounts—Google offers up to 30% off for workloads that run for most of the billing month on GCP services.

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

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

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

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

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

NEW QUESTION 136

- (Exam Topic 2)

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

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

Answer: D

Explanation:

Cloud Spanner:

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

- Get all the benefits of relational semantics and SQL with unlimited scale
- Start at any size and scale with no limits as your needs grow
- Enjoy high availability with zero scheduled downtime and online schema changes
- Deliver high-performance transactions with strong consistency across regions and continents
- Focus on innovation, eliminating manual tasks with capabilities like automatic sharding.

NEW QUESTION 141

- (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 143

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

Below are my top 5 characteristics of a DevOps organization.

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

NEW QUESTION 147

- (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 150

- (Exam Topic 2)

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

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

D. All of the above.

Answer: D

Explanation:

When you create a virtual machine (VM) instance, you must also create a boot disk for the VM. You can use a public image, a custom image, or a snapshot that was taken from another boot disk. When you create a boot disk, limit the disk size to 2 TB to account for the limitations of MBR partitioning. Compute Engine automatically creates a boot persistent disk when you create an instance. If you require additional data storage space for your instances, add one or more secondary instance storage options. You might need to create a standalone boot persistent disk and attach it to an instance later, or resize a boot persistent disk to improve performance and add more space for additional applications or operating system files. That process is described in [Add or resize a persistent disk](#). As a best practice, do not use regional persistent disks for boot disks. In a failover situation, they do not force-attach to a VM. After you create a snapshot of a boot persistent disk, you can apply data from that snapshot to new persistent disks. It is only possible to apply data from a snapshot when you first create a persistent disk. You cannot apply a snapshot to an existing persistent disk, or apply a snapshot to persistent disks that belong to a different project than that snapshot.

NEW QUESTION 155

- (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 159

- (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 162

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 167

- (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 170

- (Exam Topic 2)

What are the different storage & database services in GCP? Which is Google cloud storage and da-tabase 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 173

- (Exam Topic 2)

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

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

Answer: C

Explanation:

Firebase Crashlytics:

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

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

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

NEW QUESTION 174

- (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.

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

Answer: D

Explanation:

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 178

- (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 identify the issue quickly within the strict timeline for fixes according to the contract, and then fix it. What should you do?

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

Answer: D

Explanation:

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

-> Such fixes are unlikely to change core issues like the service itself being architected or written

sub-optimally. Though changes like browser, networking, etc. are helpful, it would be the wrong approach to first recommend that the customer upgrade all their hardware and software.

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

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

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

NEW QUESTION 181

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 184

- (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 187

- (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 188

- (Exam Topic 2)

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

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

Answer: AB

NEW QUESTION 190

- (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 195

- (Exam Topic 2)

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

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

Answer: A

Explanation:

AutoML Vision Edge model can be deployed to one of several types of edge devices, such as mobile phones, ARM-based devices, and the Coral Edge TPU

<https://cloud.google.com/vision/automl/docs/edge-quickstart>

NEW QUESTION 198

- (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 199

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 203

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 207

- (Exam Topic 2)

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

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

Answer: B

Explanation:

Because unsecured third-party systems are a cybersecurity threat.

NEW QUESTION 208

- (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 210

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 213

- (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 214

- (Exam Topic 3)

Why should an organization consider the total cost of ownership (TCO) when moving from on-premises to the cloud?

- A. To evaluate error budget
- B. To understand service level availability
- C. To evaluate return on investment
- D. To calculate required compute power

Answer: C

NEW QUESTION 216

- (Exam Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 217

- (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 218

- (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 223

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 224

- (Exam Topic 3)

An organization wants to create a new application in the cloud to replace an existing on-premises application. Which application modernization approach should the organization use?

- A. Move the application to the cloud, and then change it.
- B. Change their application, and then move it to the cloud.

- C. Invent in greenfield.
- D. Invent in brownfield.

Answer: D

Explanation:

This approach carries over as much custom components as possible from the source system and minimizes initial reengineering efforts.

NEW QUESTION 225

- (Exam Topic 3)

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 229

- (Exam Topic 3)

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

Which infrastructure solution should they choose?

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

Answer: A

Explanation:

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

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

NEW QUESTION 230

- (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 232

- (Exam Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 233

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 236

- (Exam Topic 3)

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

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

Answer: B

NEW QUESTION 241

- (Exam Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 244

- (Exam Topic 3)

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

- A. They can expect 100% service availability.
- B. They can avoid the limitations of serverless computing.
- C. They can have new tools to innovate and optimize resource usage.
- D. They can optimize maintenance for their on-premises infrastructure.

Answer: C

Explanation:

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

NEW QUESTION 249

- (Exam Topic 3)

A retail company stores their product inventory in a legacy system. Often, customers find products on the company's website and want to purchase them in-store. However, when they arrive, they discover that the products are out of stock. How could the company benefit from using an application programming interface (API)?

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

Answer: D

Explanation:

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

NEW QUESTION 254

- (Exam Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 255

- (Exam Topic 3)

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

What should the organization use?

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

Answer: B

NEW QUESTION 257

- (Exam Topic 3)

An organization needs to store structured, semi-structured, and unstructured data in its raw, native format in the same repository.

Which cloud data management solution should the organization use?

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 261

.....

Relate Links

100% Pass Your Cloud-Digital-Leader Exam with Examible Prep Materials

<https://www.exambible.com/Cloud-Digital-Leader-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>