

Google

Exam Questions Cloud-Digital-Leader

Google Cloud Digital Leader exam



NEW QUESTION 1

- (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 2

- (Topic 1)

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

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

Answer: C

Explanation:

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

Combining reservations with commitments

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

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

NEW QUESTION 3

- (Topic 1)

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

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

Answer: B

Explanation:

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

```
# Template for new folder & new project

folder_resource = {
  'name': 'new-folder',
  'type': 'gcp-types/cloudresourcemanager-v2:folders',
  'properties': {
    'parent': 'organizations/999999',
    'displayName': 'new-folder'
  }
}

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

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

Rectangular Snip

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

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

NEW QUESTION 4

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 5

- (Topic 1)

Your organization is building an application running in Google Cloud. Currently, software builds, tests, and regular deployments are done manually, but you want to reduce work for the team. Your organization wants to use Google Cloud managed solutions to automate your build, testing, and deployment process.

Which Google Cloud product or feature should your organization use?

- A. Cloud Scheduler
- B. Cloud Code
- C. Cloud Build
- D. Cloud Deployment Manager

Answer: C

Explanation:

Deploy your application to App Engine using the gcloud app deploy command. This command automatically builds a container image by using the Cloud Build service and then deploys that image to the App Engine flexible environment.

Reference: <https://cloud.google.com/appengine/docs/flexible/nodejs/testing-and-deploying-your-app>

NEW QUESTION 6

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 7

- (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 8

- (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 9

- (Topic 1)

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

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

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

Answer: C

Explanation:

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

Reference: <https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-introduction> Text Description automatically generated <https://support.google.com/a/answer/106368?hl=en>

NEW QUESTION 10

- (Topic 1)

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

What should your organization do?

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

Answer: C

Explanation:

Graphical user interface, text, application, Word, email

Direct Peering overview

[Send feedback](#)

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

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

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

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

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

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

NEW QUESTION 10

- (Topic 1)

What are the key features of Google Cloud Identity.

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

Answer: D

Explanation:

Cloud Identity:

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

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

KEY FEATURES :

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

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

Endpoint management

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

Single sign-on (SSO)

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

Works with your favorite apps

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

NEW QUESTION 15

- (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 17

- (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:

Resource hierarchy	Structure and organize your resource hierarchy for fine-grained management and cost allocation using organizations, folders, projects, and labels.
Billing access control	Enforce organizational policies with granular permissions at different levels in the resource hierarchy to control who can spend and who has administrative and cost-viewing permissions.

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 20

- (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.

- **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 22

- (Topic 1)

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

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

Answer: B

Explanation:

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

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

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

References:

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

NEW QUESTION 26

- (Topic 1)

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

- A. Customers expect faster time to market for games.
- B. Employees expect source code changes to be deployed faster.
- C. Customers expect a personalized experience.
- D. Employees expect more predictable data management spending.

Answer: C

Explanation:

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

NEW QUESTION 28

- (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 29

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 34

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 38

- (Topic 3)

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

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

D. Cloud Debugger

Answer: A

Explanation:

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

NEW QUESTION 39

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 44

- (Topic 3)

How is privacy defined in the context of cloud technology?

- A. Restrictions on data access and sharing
- B. Procedures to authenticate user identity
- C. Susceptibility to data breaches and cyber attacks
- D. Compliance with regulatory standards

Answer: A

NEW QUESTION 46

- (Topic 3)

Why is data stored in Google Cloud secure and private?

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

Answer: C

Explanation:

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

NEW QUESTION 47

- (Topic 3)

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

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

Answer: A

NEW QUESTION 48

- (Topic 3)

An organization needs to migrate specialized workloads to the cloud while maintaining their existing complex licensing and architecture. What Google Cloud solution should the organization use?

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

Answer: B

Explanation:

"This solution provides a path to modernize your application infrastructure landscape, while maintaining your existing investments and architecture. With Bare Metal Solution, you can bring your specialized workloads to Google Cloud, allowing you access and integration with GCP services with minimal latency."

NEW QUESTION 50

- (Topic 3)

An organization is struggling to keep up with the growth of their application which is running on legacy infrastructure. What might be holding them back?

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

Answer: D

Explanation:

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

NEW QUESTION 53

- (Topic 3)

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

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

Answer: C

NEW QUESTION 57

- (Topic 3)

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

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

Answer: C

NEW QUESTION 59

- (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 61

- (Topic 3)

What is an example of structured data that a healthcare facility stores in their system?

- A. X-ray images
- B. Surgery video recordings
- C. Blood pressure history
- D. Physician-written notes

Answer: C

Explanation:

Physical measures like height, weight, blood pressure, blood type, and stage of the disease can be recorded numerically and they are structured.

NEW QUESTION 62

- (Topic 3)

A global organization is developing an application to manage payments and online bank accounts in multiple regions. Each transaction must be handled consistently in their database, and they anticipate almost unlimited growth in the amount of data stored. Which Google Cloud product should the organization choose?

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

Answer:

B

NEW QUESTION 67

- (Topic 3)

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

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

Answer: C

NEW QUESTION 72

- (Topic 3)

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

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

Answer: A

NEW QUESTION 75

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 78

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 82

- (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 83

- (Topic 3)

Which policy helps Google Cloud keep customer data private?

- A. Google tests the service availability of customer applications.
- B. Google does not use customer data for advertising purposes.
- C. Google migrates customer data to an offline server when a threat is detected.
- D. Google does not allow customers to change encryption keys.

Answer: B

NEW QUESTION 87

- (Topic 3)

An organization is moving away from an on-premises infrastructure. Instead, they want to create, access, and share information virtually in the cloud. What should the organization consider?

- A. Built-in security when moving their data to the cloud
- B. Replacing their perimeter security with data encryption keys
- C. Optimizing cost-management with a capital expenditure model
- D. Increased hardware capacity when moving their data to the cloud

Answer: A

NEW QUESTION 89

- (Topic 3)

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

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

Answer: C

NEW QUESTION 94

- (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 97

- (Topic 3)

An organization recently launched a virtual customer support agent, generating vast amounts of text and speech data. Why should they use a cloud data warehouse to interpret this data?

- A. To natively visualize both types of data using a dashboard in real time
- B. To ingest and analyze structured and unstructured data at scale, in real time
- C. To secure data transmission between cloud and on-premises environments
- D. To transform data from structured to unstructured

Answer: B

Explanation:

Real-time data ingestion and updates. A simple and universal solution for continually ingesting your enterprise data into popular cloud-based data warehouses in real time. <https://www.qlik.com/us/cloud-data-migration/cloud-data-warehouse>

NEW QUESTION 102

- (Topic 3)

An organization wants full control of their virtual machine infrastructure for a custom home-grown application with a product that autoscales and automatically updates.

Which Google Cloud product or solution should the organization use?

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

Answer: C

Explanation:

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

NEW QUESTION 104

- (Topic 3)

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

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

Answer: C

NEW QUESTION 109

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 111

- (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 113

- (Topic 3)

What is logging within the context of cloud technology?

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

Answer: A

Explanation:

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

NEW QUESTION 116

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 117

- (Topic 3)

What is an organization exclusively responsible for when they access an application through a software as a service (SaaS) model?

- A. Maintaining overall system operability
- B. Maintaining customer-facing content
- C. Monitoring data center servers
- D. Monitoring computer networks

Answer: B

NEW QUESTION 120

- (Topic 3)

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

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

Answer: A

NEW QUESTION 125

- (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 128

- (Topic 3)

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 133

- (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 136

- (Topic 2)

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

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

Answer: AB

Explanation:

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

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

NEW QUESTION 139

- (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 144

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 149

- (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>

Compliance Reports Manager

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

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

Text, timeline Description automatically generated

Compliance Offerings – <https://cloud.google.com/security/compliance/offerings>

Compliance offerings

To help you with compliance and reporting, we share information, best practices, and easy access to documentation.

Our products regularly undergo independent verification of security, privacy, and compliance controls, achieving certifications against global standards to earn your trust. We're constantly working to expand our coverage.

This site contains information about Google's certifications and compliance standards it satisfies as well as general information about certain region or sector-specific regulations.

Text Description automatically generated

GDPR Resource Center – <https://cloud.google.com/security/gdpr/resource-center> At Google Cloud, we champion initiatives that prioritize and improve the security and

privacy of customer personal data, and want you, as a Google Cloud customer, to feel confident using our services in light of GDPR requirements. If you partner with Google Cloud, we will support your GDPR compliance efforts

NEW QUESTION 154

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 156

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 158

- (Topic 2)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 160

- (Topic 2)

What load balancer type is supported with Cloud Armor security policies?

- A. SSL Proxy, HTTP(S) and SSL
- B. HTTP(S) and SSL
- C. Regional SSL
- D. HTTP(S) Only

Answer: D

Explanation:

Google Cloud Armor security policies protect your application by providing Layer 7 filtering and by scrubbing incoming requests for common web attacks or other Layer 7 attributes to potentially block traffic before it reaches your load balanced backend services or backend buckets. Each security policy is made up of a set of rules that filter traffic based on conditions such as an incoming request's IP address, IP range, region code, or request headers.

-> Google Cloud Armor security policies are available only for backend services behind an external HTTP(S) load balancer. The load balancer can be in Premium Tier or Standard Tier.

-> Google Cloud Armor security policies and IP DENY lists and ALLOW lists are available only for HTTP(S) load balancing.

Reference link- <https://cloud.google.com/armor/docs/security-policy-overview>

NEW QUESTION 163

- (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 cloud
- D. Thus re-sponses also will be fast.
- E. Use TensorFlow to create custom models and deploy it as TensorFlow Lite models.

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 167

- (Topic 2)

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

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

Answer: E

Explanation:

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.

Graphical user interface, text, application Description automatically generated

NEW QUESTION 170

- (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 175

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 177

- (Topic 2)

An organization wants to measure everything as part of its new DevOps philosophy. What should the organization measure?

- A. The reliability and health of their systems.
- B. The satisfaction and happiness of their employees.
- C. The risk and reward of their investments.
- D. The speed of their cloud adoption process.

Answer: A

Explanation:

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

DevOps measurements for reliability and system health

DevOps teams can track system reliability, quality, and overall health using a few key metrics. In DevOps organizations, site reliability engineers, operations engineers, software developers, project managers, and engineering leadership will all find value in these measurements.

<https://newrelic.com/devops/measuring-devops#toc-devops-measurments-for-team-health>

NEW QUESTION 182

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 185

- (Topic 2)

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

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

Answer: D

Explanation:

What Bare Metal Solution provides

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

With Bare Metal Solution, Google Cloud provides and manages the core infrastructure, the network, the physical and network security, and hardware monitoring capabilities in an environment from which you can access all of the Google Cloud services. The core infrastructure includes secure, controlled-environment facilities, and power.

The Bare Metal Solution also includes the provisioning and maintenance of custom, sole- tenancy servers with local SAN, and smart hands support.

The network, which is managed by Google Cloud, includes a low-latency Partner Interconnect connection into the customer Bare Metal Solution environment.

The available Google Cloud services include private API access, management tools, support, and billing.

NEW QUESTION 190

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 195

- (Topic 2)

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 200

- (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 202

- (Topic 2)

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

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 203

- (Topic 2)

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

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

Answer: AD

Explanation:

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

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

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

NEW QUESTION 205

- (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 anyway 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 207

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 210

- (Topic 2)

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

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

Answer: D

Explanation:

Cloud Spanner:

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

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

NEW QUESTION 211

- (Topic 2)

Which of the following statements is/are true about Google Cloud BigTable?

- A. It is not compatible with Hadoop.
- B. It Scales from Giga Byte to Peta Byte with No Downtime.
- C. It can not be used in Real-time Ad analytics and tracking thousands of IoT Devices Data.
- D. It is an enterprise-level Database that offers relational and non-relational features

Answer: B

Explanation:

Cloud Bigtable

A fully managed, scalable NoSQL database service for large analytical and operational workloads with up to 99.999% availability.

- Consistent sub-10ms latency—handle millions of requests per second
- Ideal for use cases such as personalization, ad tech, fintech, digital media, and IoT
- Seamlessly scale to match your storage needs; no downtime during reconfiguration
- Designed with a storage engine for machine learning applications leading to better predictions
- Easily connect to Google Cloud services such as BigQuery or the Apache ecosystem

NEW QUESTION 215

- (Topic 2)

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

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

Answer: D

Explanation:

Cloud SQL

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

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

Key features:

Fully managed

Cloud SQL automatically ensures your databases are reliable, secure, and scalable so that your business continues to run without disruption. Cloud SQL automates all your backups, replication, encryption patches, and capacity increases—while ensuring greater than 99.95% availability, anywhere in the world.

Integrated

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

Easily configure replication and backups to protect your data. Go further by enabling automatic failover to make your database highly available. Your data is

automatically encrypted, and Cloud SQL is SSAE 16, ISO 27001, and PCI DSS compliant and supports HIPAA compliance.

Easy migrations to Cloud SQL

Database Migration Service (DMS) makes it easy to migrate your production databases to Cloud SQL with minimal downtime. This serverless offering eliminates the manual hassle of provisioning, managing, and monitoring migration-specific resources. DMS leverages the native replication capabilities of MySQL and PostgreSQL to maximize the fidelity and reliability of your migration. And it's available at no additional charge for native like-to-like migrations to Cloud SQL.

NEW QUESTION 216

- (Topic 2)

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

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

Answer: B

Explanation:

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

Instance groups

[Send feedback](#)

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

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

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

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

NEW QUESTION 217

- (Topic 2)

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

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

Answer: C

Explanation:

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

Cloud SQL Key Terms: Cloud SQL instance

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

Database instance

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

High availability

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

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

Failover

A failover is when Cloud SQL switches serving from the original primary instance to the standby instance.

Autofailover is a mechanism that automatically triggers failover when a Cloud SQL instance didn't issue a heartbeat in the previous interval.

Standby instances

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

Clone

When you clone a Cloud SQL instance, you create a new instance that is a copy of the source instance, but is completely independent. After cloning is complete, changes to the

source instance are not reflected in the clone, and changes in the clone are not reflected in the source instance.

Replication

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

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

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

Source server

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

Cloud SQL Auth proxy client

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

NEW QUESTION 221

- (Topic 2)

You are a database manager working for a new product that will need millions of reading and writing from the database, with zero downtime, key-value i.e. NoSQL features, no manual steps should be required to ensure consistency, repair data, synchronize writes and deletes, Which of the following database you choose?

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

Answer: B

Explanation:

Cloud BigTable

Key features

High throughput at low latency

Bigtable is ideal for storing very large amounts of data in a key-value store and supports high read and write throughput at low latency for fast access to large amounts of data. Throughput scales linearly—you can increase QPS (queries per second) by adding Bigtable nodes. Bigtable is built with proven infrastructure that powers Google products used by billions such as Search and Maps.

Cluster resizing without downtime

Scale seamlessly from thousands to millions of reads/writes per second. Bigtable throughput can be dynamically adjusted by adding or removing cluster nodes without restarting, meaning you can increase the size of a Bigtable cluster for a few hours to handle a large load, then reduce the cluster's size again—all without any downtime. Flexible, automated replication to optimize any workload

Write data once and automatically replicate where needed with eventual consistency—giving you control for high availability and isolation of reading and write workloads. No manual steps are needed to ensure consistency, repair data, or synchronize writes and deletes. Benefit from a high availability SLA of 99.999% for instances with multi-cluster routing across 3 or more regions (99.9% for single-cluster instances).

NEW QUESTION 225

- (Topic 2)

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

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 228

- (Topic 2)

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

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

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

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

Answer: C

Explanation:

The above is a real world example of a company named Planet, where they sent around 80+ satellites to take pictures of earth every day, 24 hours. They run around 40,000 preemptible VMs concurrently. Preemptible instances function like normal instances but have the following limitations: Compute Engine might stop preemptible instances at any time due to system events. The probability that Compute Engine will stop a preemptible instance for a system event is generally low, but might vary from day to day and from zone to zone depending on current conditions. Compute Engine always stops preemptible instances after they run for 24 hours. Certain actions reset this 24-hour counter. Preemptible instances are finite Compute Engine resources, so they might not always be available. Preemptible instances can't live migrate to a regular VM instance, or be set to automatically restart when there is a maintenance event. Due to the above limitations, preemptible instances are not covered by any Service Level Agreement (and, for clarity, are excluded from the Compute Engine SLA). The Google Cloud Free Tier credits for Compute Engine do not apply to preemptible instances.

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

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

NEW QUESTION 231

- (Topic 2)

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

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

Answer: B

Explanation:

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

interface, text, application, email Description automatically generated

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

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

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

NEW QUESTION 233

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 237

- (Topic 2)

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

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

Answer: B

Explanation:

Cloud Spanner:

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

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

Automatic sharding

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

Strong transactional consistency

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

Regional and multi-regional configurations

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

Online schema changes with no downtime

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

NEW QUESTION 241

- (Topic 1)

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

What should you do?

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

Answer: C

Explanation:

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

NEW QUESTION 246

- (Topic 1)

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

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

Answer: A

Explanation:

Reference: <https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets>

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

<https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets>

NEW QUESTION 247

- (Topic 1)

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

- A. Cloud Functions
- B. App Engine
- C. Compute Engine
- D. Anthos

Answer: D

Explanation:

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

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

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

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

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

NEW QUESTION 251

- (Topic 1)

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

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

Answer: A

Explanation:

A VM interface can send packets to the external IP addresses of Google APIs and services using Private Google Access if all these conditions are met:

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

NEW QUESTION 256

- (Topic 1)

What are the network requirements for Private Google Access?

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

Answer: B

Explanation:

Network requirements for Private Google Access:

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

NEW QUESTION 259

- (Topic 1)

You are a program manager within a Software as a Service (SaaS) company that offers rendering software for animation studios. Your team needs the ability to allow scenes to be scheduled at will and to be interrupted at any time to restart later. Any individual scene rendering takes less than 12 hours to complete, and there is no service-level agreement (SLA) for the completion time for all scenes. Results will be stored in a global Cloud Storage bucket. The compute resources are not bound to any single geographical location. This software needs to run on Google Cloud in a cost-optimized way. What should you do?

- A. Deploy the application on Compute Engine using preemptible instances

- B. Develop the application so it can run in an unmanaged instance group
- C. Create a reservation for the minimum number of Compute Engine instances you will use
- D. Start more instances with fewer virtual centralized processing units (vCPUs) instead of fewer instances with more vCPUs

Answer: A

Explanation:

What is a preemptible instance?

Preemptible VM instances are available at much lower price—a 60-91% discount—compared to the price of standard VMs. However, Compute Engine might stop (preempt) these instances if it needs to reclaim the compute capacity for allocation to other VMs. Preemptible instances use excess Compute Engine capacity, so their availability varies with usage.

If your apps are fault-tolerant and can withstand possible instance preemptions, then preemptible instances can reduce your Compute Engine costs significantly. For example, batch processing jobs can run on preemptible instances. If some of those instances stop during processing, the job slows but does not completely stop. Preemptible instances complete your batch processing tasks without placing additional workload on your existing instances and without requiring you to pay full price for additional normal instances.

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

NEW QUESTION 260

- (Topic 1)

Your organization needs to analyze data in order to gather insights into its daily operations. You only want to pay for the data you store and the queries you perform. Which Google Cloud product should your organization choose for its data analytics warehouse?

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

Answer: D

Explanation:

BigQuery is an enterprise data warehouse for large amounts of relational structured data Serverless, highly scalable, and cost-effective multicloud data warehouse designed for business agility.

NEW QUESTION 263

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 264

- (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 268

- (Topic 1)

There are internal compliance requirements that demand that we do not use any APIs or services that are not backed by SLAs. Which of these are acceptable for

us? (Choose two answer)

- A. Alpha, Beta
- B. Early Access, Preview
- C. General Availability
- D. Deprecated, but ensure that the SLA support period is still valid.

Answer: CD

Explanation:

General Availability is the stage where SLAs apply.

Deprecated - in the deprecated stage, you should start moving away from those APIs and products. Depending on the deprecation policy, SLAs could still be valid.

NEW QUESTION 271

- (Topic 1)

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

- A. The network, which Google Cloud manages includes a low-latency Cloud Inter-connect connection into the customer Bare Metal Solution environment.
- B. Bare Metal Solution also includes the provisioning and maintenance of the cus-tom, sole-tenancy hardware with local SAN, and smart hands support.
- C. Bare Metal Solution uses a bring-your-own-license (BYOL) model.
- D. All of the Above.

Answer: D

Explanation:

Option A is true

You are responsible for the licensing of all of your software. Bare Metal Solution uses a bring-your-own-license (BYOL) model.

Apart from this you are responsible for the software, applications, and data that you use and store in the Bare Metal Solution environment.

Responsibilities Data, including:

- Security and encryption
- Backups

Software and applications, including:

- Installation
- Configuration
- Upgrades and patching

Operating system and any hypervisor, including:

- Configuration changes
- Upgrades and patching Server clusters, including:
 - Installation
 - Configuration
 - Maintenance Licensing

Option B & C is also true.

With Bare Metal Solution, Google Cloud provides and manages the core infrastructure, the net-work, the physical and network security, and hardware monitoring capabilities in an

environment from which you can access all of the Google Cloud services. The core infrastructure includes secure, controlled-environment facilities, and power.

The Bare Metal Solution also includes the provisioning and maintenance of the custom, sole-tenancy hardware with local SAN, and smart hands support.

The network, which is managed by Google Cloud includes a low-latency Cloud Interconnect con-nection into the customer Bare Metal Solution environment.

The available Google Cloud services include private API access, management tools, support, and billing.

NEW QUESTION 272

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

Cloud-Digital-Leader Practice Exam Features:

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

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