Cloud-Digital-Leader Dumps

Google Cloud Digital Leader exam

https://www.certleader.com/Cloud-Digital-Leader-dumps.html

NEW QUESTION 1
- (Topic 1)
You are leading projects in an IT services company. Your customer's project requires analyzing im-ages. 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 needs to minimize how much it pays for data traffic from the Google network to the internet. What should your organization do?

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

Answer: A

Explanation:

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

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

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

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

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

NEW QUESTION 3
- (Topic 1)
Your application has repeated data requests of the exact same nature. At the same time, the number of user requests is increasing. Monitoring indicates that the load on the existing database is increasing, and there seems to be a bottleneck. An analysis of the data requested shows us that it is application-managed data and that it changes, but not often. How can you improve the efficiency of the application?

A. Use Cloud Memorystore to improve speed via caching
B. Increase the amount of RAM on the machine hosting the database so that it has higher data throughput.
C. Use Cloud Storage with multi-regional storage so that all users accessing the data will have lower latency
D. Increase the number of CPUs on the machine hosting the database so that it has higher data throughput.

Answer: A

Explanation:
Cloud Memorystore is an in-memory database that has sub-millisecond latency. This is ideal for caching application data that also changes once in a while.
https://cloud.google.com/memorystore

NEW QUESTION 4
- (Topic 1)
Which Google Cloud product can report on and maintain compliance on your entire Google Cloud organization to cover multiple projects?

A. Cloud Logging

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

Answer: D

Explanation:
 Security Command Center is a centralized security and risk management platform for your Google Cloud resources. It is a single tool that offers a variety of security features including:
* 1. Gain centralized visibility and control
* 2. Discover misconfigurations and vulnerabilities
* 3. Report on and maintain compliance
* 4. Detect threats targeting your Google Cloud assets https://cloud.google.com/security-command-center

NEW QUESTION 5
- (Topic 1)
Each of the three cloud service models - infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS) - offers benefits between flexibility and levels of management by the cloud provider and the customer.
Why would SaaS be the right choice of service model?

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

Answer: B

Explanation:
 Benefits of SaaS
The main benefit of SaaS is that it offloads all infrastructure and application management to the SaaS vendor.
Reference: https://www.ibm.com/cloud/learn/iaas-paas-saas

NEW QUESTION 6
- (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 7
- (Topic 1)
Your organization is developing a plan for migrating to Google Cloud.
What is a best practice when initially configuring your Google Cloud environment?

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

Answer: B

Explanation:
 The Organization resource is the root node of the Google Cloud resource hierarchy and all resources that belong to an organization are grouped under the organization node. This provides central visibility and control over every resource that belongs to an organization.
Reference link- https://cloud.google.com/resource-manager/docs/cloud-platform-resource- hierarchy

NEW QUESTION 8
- (Topic 1)
Your organization wants to predict the behavior of visitors to its public website. To do that, you have decided to build a machine learning model. Your team has database-related skills but only basic machine learning skills, and would like to use those database skills.
Which Google Cloud product or feature should your organization choose?

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

Answer: A

Explanation:

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

NEW QUESTION 9
- (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 10
- (Topic 1)
Your organization needs to process large amounts of data from an online application that operates continuously. You do not want to be required to provision infrastructure or create server clusters. What should your organization choose?

A. Compute Engine with BigQuery
B. Dataproc
C. Google Kubernetes Engine with Cloud Bigtable
D. Dataflow

Answer: D

Explanation:
 You do not want to be required to provision infrastructure or create server clusters. Because Unified stream and batch data processing that's serverless, fast, and cost-effective.
Reference link- https://cloud.google.com/dataflow

NEW QUESTION 10
- (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 12
- (Topic 1)
Which Google Cloud product is designed to reduce the risks of handling personally identifiable information (PII)?

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

Answer: C

Explanation:
Reference: https://cloud.google.com/blog/products/gcp/take-charge-of-your-sensitive-data-with-the-cloud-dlp-api
Cloud Data Loss Prevention: Fully managed service designed to help you discover, classify, and protect your most sensitive data.

NEW QUESTION 14
- (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 18
- (Topic 1)
Which Google Cloud service or feature lets you build machine learning models using Standard SQL and data in a data warehouse?

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

Answer: A

Explanation:
BigQuery ML lets you create and execute machine learning models in BigQuery using standard SQL queries.
Reference: https://cloud.google.com/bigquery-ml/docs/introduction#:~:text=BigQuery%20ML%20lets%20you%20create,the%20need%20to%20move%20data
Graphical user interface, text, application, email Description automatically generated
https://cloud.google.com/bigquery-ml/docs/introduction

NEW QUESTION 22
- (Topic 1)
Your team has developed a machine learning model for your customer. The test results indicate very strong predictive capability. The model is then deployed in production. Evaluation of the predictions in production show that they are off by a pronounced margin. What is the issue and how can you solve for it?

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

Answer: D

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

NEW QUESTION 26
- (Topic 3)
An organization is looking for a storage solution that will help them serve content to users worldwide. They need a solution that offers a high level of availability What feature of Cloud Storage would they benefit from?

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

Answer: D

NEW QUESTION 30
- (Topic 3)
A large retail organization uses traditional technology for their ecommerce website During peaks m traffic resources are often underutilized or overprovisioned They have decided to migrate to cloud technology
What aspect of cloud technology will benefit their ecommerce business?

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

Answer: A

NEW QUESTION 34
- (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 37
- (Topic 3)
An organization's web developers and operations personnel use different systems. How will increasing communication between the teams reduce issues caused by silos?

A. By assigning blame for failures and establishing consequences
B. By combining job role responsibilities to ensure that everyone has shared access
C. By increasing data encryption to strengthen workflows
D. By emphasizing shared ownership of business outcomes

Answer: D

NEW QUESTION 42
- (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 43
- (Topic 3)
An organization is making a strategic change to customer support in response to feedback. They plan to extend their helpline availability hours.
Why is the organization making this change?

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

Answer: C

NEW QUESTION 47
- (Topic 3)
How is service availability measured in the context of cloud technology?

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

Answer: B

NEW QUESTION 48
- (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)
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 55
- (Topic 3)
An online retail organization wants to optimize their service.
What is an example of unstructured data that they can use to make decisions?

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

Answer: A

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

NEW QUESTION 58
- (Topic 3)
Why do organizations often struggle to scale their on-premises application infrastructure?

A. Scaling compute instances could breach compliance and/or regulation
B. Increasing compute capacity is time-consuming and costly
C. Their serverless compute functions struggle to meet the demand
D. Their multi-cloud architecture is complex and expensive

Answer: B

NEW QUESTION 61
- (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 62
- (Topic 3)
An organization has an on-premises IT infrastructure. Their customer-facing application repeatedly fails during peak usage.
What could be causing this issue?

A. A serverless compute function struggles to scale.
B. The application contains unclean data.
C. They don't have enough servers to meet the demand.
D. The application is only configurable on-premises.

Answer: C

NEW QUESTION 64
- (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 66
- (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 70
- (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 72
- (Topic 3)
An organization wants to use multiple marketing datasets to forecast user acquisition. How should they use cloud technology to gain new insights from the data?

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

Answer: D

NEW QUESTION 73
- (Topic 3)
An organization wants to develop an application that can be personalized to user preferences throughout the year.
Why should they build a cloud-native application instead of modernizing their existing on- premises application?

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

Answer: B

NEW QUESTION 74
- (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 79
- (Topic 2)
Your client is building a custom machine learning pipeline to identify lesions in the lungs based on x-rays. Different teams of data scientists are sharing common source data and building many ver-sions of ML models. Which of these Cloud Storage options would be best for them?

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

Answer: B

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

Standard storage

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

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

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

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

Availability

The availability of Standard storage data is:

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

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

NEW QUESTION 81
- (Topic 2)
You are discussing scaling requirements with a gaming company. When the game launches, they are expecting incoming data surges of 2 million users or more during weekends and holidays. Their on-premise systems have had issues scaling and they want your advice on solving the issue. What do you recommend?

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

Answer: B

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

NEW QUESTION 84
- (Topic 2)
A customer deploys an application to App Engine and needs to check for Open Web Appli- cation Security Project (OWASP) vulnerabilities. Which service should be used to accom- plish this?

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

Answer: B

Explanation:
Web Security Scanner identifies security vulnerabilities in your App Engine, Google Kubernetes Engine (GKE), and Compute Engine web applications. It crawls your application, following all links within the scope of your starting URLs, and attempts to exercise as many user inputs and event handlers as possible. Currently, Web Security Scanner only supports public URLs and IPs that aren't behind a firewall. Web Security Scanner currently supports the App Engine standard environment and App Engine flexible environments, Compute Engine instances, and GKE resources.
Reference link- https://cloud.google.com/security-command-center/docs/concepts-web- security-scanner-overview

NEW QUESTION 88
- (Topic 2)
Your company has made plans to roll out OpenShift, a Kubernetes platform solution offered by IBM Red Hat, across all its on-premises and public cloud environments. Given that you are the lead architect responsible for your company's GCP deployments, what type of shared responsibility model will this deployment entail for you?

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

Answer: D

Explanation:
The key to remember here is that for a service provided (GCP in this case) to take responsibility for its PaaS, it must offer the service as a managed service. GCP

offers its own Kubernetes platform called GKE. But OpenShift is not a Google-offered PaaS solution. As such, Google will not take responsibility for the back-end operations and design of your OpenShift environments. You will need to manage all the VMs that OpenShift will provision as part of its GCP deployment. So this is an IaaS deployment from a shared responsibility model perspective.

NEW QUESTION 93
- (Topic 2)
You are storing sensitive information in a Cloud Storage bucket. For legal reasons, you need to be able to record all requests that read any of the stored data. You want to make sure you comply with these requirements. What should you do?

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

Answer: B

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

Available audit logs

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

- Admin Activity audit logs: Entries for `ADMIN_WRITE` operations that modify the configuration or metadata of a Cloud project, bucket, or object. You can't disable Admin Activity audit logs.

- Data Access audit logs: Entries for operations that modify objects or read a Cloud project, bucket, or object. There are several sub-types of Data Access audit logs:

  - `ADMIN_READ` : Entries for operations that read the configuration or metadata of a Cloud project, bucket, or object.

  - `DATA_READ` : Entries for operations that read an object.

  - `DATA_WRITE` : Entries for operations that create or modify an object.

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

For fuller descriptions of the audit log types, see Types of audit logs.

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

NEW QUESTION 97
- (Topic 2)
A large organization is struggling to manage their cloud costs effectively. They want to increase vis-ibility into cloud costs. Which cost management approach should the organization use?

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

Answer: A

Explanation:
 Because cross-team partnerships are part of the visibility cost management strategy.
https://wa.aws.amazon.com/wat.question.COST_1.en.html

NEW QUESTION 99
- (Topic 2)
Google offers Firebase, In terms of Firebase Console, any particular message that has to be deliv-ered to a customer at a certain degree of change in behavior can be managed through_____.

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

Answer: B

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

NEW QUESTION 103

- (Topic 2)
certain devices for cracks, rust, etc. Some of these issues are difficult to identify for a human and your company has seen increasing customer complaints - the customer has paid for an inspection and the field agent said there was no problem, but it later turned out there actually was. The team has come up with a proposal to engage AI to identify issues. On evaluating the existing system, it is seen that the mobile phone network connection is not good or consistent. What solution can work for them?

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

Answer: A

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

NEW QUESTION 104
- (Topic 2)
Your client has an on-premises data center. Due to technical limitations, they are unable to scale globally. They have decided to adopt the public cloud. However, they don't want to locked into any one vendor and, therefore, would like to work with multiple cloud providers. They have used open source container technologies and would like to continue using them.

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

Answer: D

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

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

NEW QUESTION 106
- (Topic 2)
In terms of Dockers and Kubernetes, which of the following statements are correct?

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

Answer: D

Explanation:
Kubernetes vs. Docker
Often misunderstood as a choice between one or the other, Kubernetes and Docker are different yet complementary technologies for running containerized applications.
Docker lets you put everything you need to run your application into a box that can be stored and opened when and where it is required. Once you start boxing up your applications, you need a way to manage them; and that's what Kubernetes does. Kubernetes is a Greek word meaning 'captain' in English. Like the captain is responsible for the safe journey of the ship in the seas, Kubernetes is responsible for carrying and delivering those boxes safely to locations where they can be used.
- Kubernetes can be used with or without Docker.
- Docker is not an alternative to Kubernetes, so it's less of a "Kubernetes vs. Docker" question. It's about using Kubernetes with Docker to containerize your applications and run them at scale.
- The difference between Docker and Kubernetes relates to the role each play in containerizing and running your applications.
- Docker is an open industry standard for packaging and distributing applications in containers.
- Kubernetes uses Docker to deploy, manage, and scale containerized applications.

NEW QUESTION 107
- (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 108
- (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 110
- (Topic 2)
What type of cloud computing service provides raw compute, storage, and network, organized in ways that are familiar to physical data centers?

A. Database as a Service.
B. Platform as a Service.
C. Infrastructure as a Service.
D. Software as a Service.

Answer: C

Explanation:
What is Infrastructure as a service :
IaaS (infrastructure as a service) is a computing model that offers resources on-demand to businesses and individuals via the cloud.
IaaS is attractive because acquiring computing resources to run applications or store data the traditional way requires time and capital. Enterprises must purchase equipment through procurement processes that can take months. They must invest in physical spaces: typically specialized rooms with power and cooling. And after deploying the systems, enterprises need IT, professionals, to manage them.
All this is challenging to scale when demand spikes or the business grows. Enterprises risk running out of capacity or overbuilding and ending up with infrastructure that suffers from low utilization.
These challenges are why IaaS use is steadily growing. Learn more about Compute Engine, Cloud Storage, etc.

NEW QUESTION 111
- (Topic 2)
Your Coustomer's Organization has decided to move to the cloud. They currently run VMs on-promise but their goal on Google cloud is to run containers, primarily on Google Kuber- nete's Engine. They have a lease for their private data center for another year that they have already paid for. What could be strategy they could adopt in migrating?

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

Answer: B

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

NEW QUESTION 115
- (Topic 2)
You are a program manager in a company and handling a project and you need to create a virtual machine on google cloud console that will be very simple to set up, by flipping a bit via command, API, or with developer console that gives you 30 seconds to shut down when you're preempted, allow you to save your work that also helps in the company budget upto 70-80% of less charges than the regular VMs.

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

Answer: B

Explanation:

Preemptible VMs have all these features
Simple configuration
Create a preemptible instance simply by flipping a bit via command, API, or developer console.
Easy extensibility
Attach GPUs and local SSDs to preemptible instances for additional performance and savings.
Graceful shutdown
Compute Engine gives you 30 seconds to shut down when you're preempted, letting you save your work in progress for later.
Large scale computing
Spin up as many instances as you need and turn them off when you're done. You only pay for what you use.
Quickly reclaim capacity
Managed instance groups automatically recreate your instances when they're preempted (if capacity is available).
Fixed pricing
Preemptible VMs have fixed pricing up to 80% off regular instances. They show up on your bill separately so you'll see just how much you're saving.

NEW QUESTION 118
- (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 121
- (Topic 2)
A large travel services company has been running all their workloads on Google Cloud in the previous year. They looked at their past usage of cloud resources and see that there is a consistent use of 10,000 virtual machines throughout the year. Based on the projections for the following year they have a strong indication that they will use at least this much or more capacity within Google Cloud. What is one way in which they can take advantage of this knowledge?

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

Answer: C

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

NEW QUESTION 124
- (Topic 2)
"With cloud messaging you can Customize and deliver messages accordingly to the predetermined time in the user's local time zone." Comment on the above statement.

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

Answer: D

Explanation:
Firebase Cloud Messaging:
Firebase Cloud Messaging (FCM) is a cross-platform messaging solution that lets you reliably send messages at no cost.
Using FCM, you can notify a client app that new email or other data is available to sync. You can send notification messages to drive user re-engagement and retention. For use cases such as instant messaging, a message can transfer a payload of up to 4000 bytes to a client app.
Key capabilities of Firebase Cloud Messaging:
Send notification messages or data messages: Send notification messages that are displayed to your user. Or send data messages and determine completely what happens in your application code.
Versatile message targeting: Distribute messages to your client app in any of 3 ways—to single devices, to groups of devices, or to devices subscribed to topics.
Send messages from client apps: Send acknowledgments, chats, and other messages from devices back to your server over FCM's reliable and battery-efficient connection channel.

NEW QUESTION 126

- (Topic 2)
What cloud deployment model is generally deployed between organizations such as non- profits, hospitals or even enterprises that share similar requirements or interests?

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

Answer: B

Explanation:
Community Cloud – The cloud infrastructure is planned for selective use by a particular community of consumers from organizations that have mutual interests like security needs, policy, and compliance considerations.
Reference link- https://csrc.nist.gov/glossary/term/community_cloud

NEW QUESTION 128
- (Topic 2)
You are looking for a one stop reference page for GCP support. What Page would you se- lect?

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

Answer: C

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

NEW QUESTION 133
- (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 136
- (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 140
- (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 144
- (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 pro-cessing. 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 147
- (Topic 1)
A fitness band company is continuously ingesting data from millions of its consumers. Different kinds of data based on time, like location, heartbeat rate, temperature, movement, etc. are connect-ed. They need a high throughput database that can write data very fast. Since their users are spread across the world, they need the database to be geographically scalable. Consumers also want to see near-real-time visualizations of their activities. Which of these databases would be a good fit?

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

Answer: B

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

NEW QUESTION 150
- (Topic 1)
An organization wants to search for and share plug-and-play AI components which can easily build ML services into their project. Which Google Cloud product should the organi- zation use?

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

Answer: B

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

NEW QUESTION 155
- (Topic 1)
Your team is using BigQuery as your central data warehouse. You are running a certain workload that you've run frequently over the last few days. It is a short, high capacity ana- lytics workload. Which of the following would be an appropriate pricing model to use?

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

Answer: C

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

NEW QUESTION 159
- (Topic 1)
You are currently managing workloads running on Windows Server for which your company owns the licenses. Your workloads are only needed during working hours, which allows you to shut down the instances during the weekend. Your Windows Server licenses are up for renewal in a month, and you want to optimize your license cost.
What should you do?

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

Answer: D

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

NEW QUESTION 164
- (Topic 1)
Your organization is developing and deploying an application on Google Cloud. Tracking your Google Cloud spending needs to stay as simple as possible.
What should you do to ensure that workloads in the development environment are fully isolated from production workloads?

A. Apply a unique tag to development resources
B. Associate the development resources with their own network
C. Associate the development resources with their own billing account
D. Put the development resources in their own project

Answer: D

Explanation:

Reference: https://www.deps.co/blog/google-cloud-platform-good-bad-ugly/
Project resources are components that are necessary for successful project implementation. They include people, equipment, money, time, knowledge – basically, anything that you may require from the project planning to the project delivery phases.

NEW QUESTION 166
- (Topic 1)
A Multiple projects within your organization have long-term VM usage. Based on current demand, they are able to project into the future and estimate how many VM hours they will use every year. Going in for a committed use contract seems sensible. How can you configure committed use easily across all the projects?

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

Answer: A

Explanation:

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

NEW QUESTION 168
......

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your Cloud-Digital-Leader Exam with Our Prep Materials Via below:

https://www.certleader.com/Cloud-Digital-Leader-dumps.html