Exam Questions Cloud-Digital-Leader

Google Cloud Digital Leader exam

https://www.2passeasy.com/dumps/Cloud-Digital-Leader/

NEW QUESTION 1
- (Topic 1)
An organization runs their application on a virtual machine, but every time they want to edit specific features, they have to bring the system offline to update the applica-tion. What would be a more appropriate solution for their app?

A. GPUs
B. Containers
C. Hypervisors
D. Solid State Disk

Answer: B

Explanation:
Because containers can compartmentalize applications which enables parts to be edited in isolation.

What are containers?

Containers are packages of software that contain all of the necessary elements to run in any environment. In this way, containers virtualize the operating system and run anywhere, from a private data center to the public cloud or even on a developer's personal laptop. From Gmail to YouTube to Search, everything at Google runs in containers. Containerization allows our development teams to move fast, deploy software efficiently, and operate at an unprecedented scale. We've learned a lot about running containerized workloads and we've shared this knowledge with the community along the way: from the early days of contributing cgroups to the Linux kernel, to taking designs from our internal tools and open sourcing them as the Kubernetes project.

Reference link- https://cloud.google.com/learn/what-are-containers

NEW QUESTION 2
- (Topic 1)
Your organization wants an economical solution to store data such as files, graphical images, and videos and to access and share them securely.
Which Google Cloud product or service should your organization use?

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

Answer: A

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

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

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

NEW QUESTION 3
- (Topic 1)
Your company has multiple internal applications used by your employees. You also have to give access to certain vendors and contractors. What is a good option for you to adopt?

A. Keep the credentials separate for each application to reduce the blast radius in case of any issues.
B. Use an external identity provider that is famous and popular like Facebook or Twitter; that way, even your vendors and contractors will have an account there.
C. Allow all users, especially contractors and vendors, to bring their own identities, like those at gmail.com.
D. Use an IDaaS (Identity as a Service) product that can centrally manage authenti-cation and authorization for the applications.

Answer: D

Explanation:

IDaaS - identity providers managed by the company give better control over security and privacy. Security/access can be set granularly, while also being centralized. You don't have to manage multiple credentials.

NEW QUESTION 4
- (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 5
- (Topic 1)
Your organization needs to allow a production job to have access to a BigQuery dataset. The production job is running on a Compute Engine instance that is part of an instance group.
What should be included in the IAM Policy on the BigQuery dataset?

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

Answer: C

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

NEW QUESTION 6
- (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/99999',
      'displayName': 'new-folder'
  }
}

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

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

Reference link- https://cloud.google.com/resource-manager/docs/creating-managing- folders
Reference link- https://stackoverflow.com/questions/59460623/how-to-create-a-folder-a-project-under-it-with-deployment-manager-google-cloud

NEW QUESTION 7
- (Topic 1)
Your organization runs all its workloads on Compute Engine virtual machine instances. Your organization has a security requirement: the virtual machines are not allowed to access the public internet. The workloads running on those virtual machines need to access BigQuery and Cloud Storage, using their publicly accessible interfaces, without violating the security requirement.
Which Google Cloud product or feature should your organization use?

A. Identity-Aware Proxy
B. Cloud NAT (network address translation)
C. VPC internal load balancers
D. Private Google Access

Answer: D

Explanation:
 VM instances that only have internal IP addresses (no external IP addresses) can use Private Google Access. They can reach the external IP addresses of Google APIs and services. The source IP address of the packet can be the primary internal IP address of the network interface or an address in an alias IP range that is assigned to the interface. If you disable Private Google Access, the VM instances can no longer reach Google APIs and services; they can only send traffic within the VPC network.

Configuring Private Google Access 🔖

<div align="right">Send feedback</div>

By default, when a Compute Engine VM lacks an external IP address assigned to its network interface, it can only send packets to other internal IP address destinations. You can allow these VMs to connect to the set of external IP addresses used by Google APIs and services by enabling Private Google Access on the subnet used by the VM's network interface.

Private Google Access also allows access to the external IP addresses used by App Engine, including third-party App Engine-based services.

To view the eligible APIs and services that you can use with Private Google Access, see supported services in the Private Google Access overview.

See Private Access Options for Services for background information about Private Google Access and other private connectivity options offered by Google Cloud.

Specifications

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 one of the following IP addresses.

   If you're sending packets to the default domains:

   - The VM interface's primary internal IPv4 address
   - The VM interface's internal IPv6 address
   - An internal IPv4 address from an alias IP range

https://cloud.google.com/vpc/docs/configure-private-google-access

NEW QUESTION 8
- (Topic 1)
Your organization is releasing its first publicly available application in Google Cloud. The application is critical to your business and customers and requires a 2-hour SLA.
How should your organization set up support to minimize costs?

A. Enroll in Premium Support
B. Enroll in Enhanced Support
C. Enroll in Standard Support
D. Enroll in Basic Support

Answer: B

Explanation:

Reference: https://www.secureauth.com/enhanced-support-offering/
SecureAuth is dedicated to providing the industry-leading enhancedsupport ensuring the long term success of your SecureAuth SaaS IAM deployment

NEW QUESTION 9
- (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 10
- (Topic 1)
You want to build an application that will allow customers to register and login. It would be great to have the ability to secure it with multi-factor authentication and the ability to reset credentials. As a small startup, you want to build the main application as quickly as possible and have minimum overhead. Which might be a suitable option for you on Google Cloud?

A. Since identity and credentials should be secure and private, do not trust other ser-vice providers.
B. Cloud Identity
C. Google Workspace
D. Cloud Identity Platform

Answer: D

Explanation:

Cloud Identity Platform
Cloud Identity Platform allows you to manage identity and credentials for your consumer- facing applications. So that's the right one in this case to use. "Identity Platform is a customer identity and access management (CIAM) platform that helps organizations add identity and access management functionality to their applications, protect user accounts, and scale with confidence on Google Cloud."
Reference link- https://cloud.google.com/identity-platform

NEW QUESTION 10
- (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 13
- (Topic 1)
An IoT platform is providing services to home security systems. They have more than a million customers, each with many home devices. Burglaries or child safety issues are concerns that the clients customers. Therefore, the platform has to respond very quickly in near real time. What could be a typical data pipeline used to support this platform on Google Cloud?

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

Answer: A

Explanation:
 Explanation
=> Cloud Pub/Sub- Cloud Pub/Sub is the best to be the end-point for ingesting large amounts of data. It will grow as required, can stream data to downstream systems, and can also work with intermittently available backends.
=> Cloud Dataflow- supports streaming data and therefore is an appropriate option for processing the data that is ingested.
=> BigQuery- BigQuery also supports streaming data and its possible to do real time ana- lytics on it.
=> DataStudio- DataStudio and Looker are for visualization. They don't have any in-built analysis.
=> Cloud Functions- Cloud Functions is a useful serverless endpoint. However, Pub/Sub is better in this case because it can also retain messages for a set period if it was not possi- ble to deliver it first time.

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

NEW QUESTION 14
- (Topic 1)
Your organization is running all its workloads in a private cloud on top of a hypervisor. Your
organization has decided it wants to move to Google Cloud as quickly as possible. Your organization wants minimal changes to the current environment, while using the maximim amount of managed services Google offers.
What should your organization do?

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

Answer: B

Explanation:
Migrate for Compute Engine enables you to lift and shift workloads at scale to Google Cloud Compute Engine with minimal changes and risk.
Reference: https://dataintegration.info/simplify-vm-migrations-with-migrate-for-compute-engine-as-a-service

NEW QUESTION 17
- (Topic 1)
You are a program manager for a team of developers who are building an event-driven application to allow users to follow one another's activities in the app. Each time a user adds himself as a follower of another user, a write occurs in the real-time database.
The developers will develop a lightweight piece of code that can respond to database writes and generate a notification to let the appropriate users know that they have gained new followers. The code should integrate with other cloud services such as Pub/Sub, Firebase, and Cloud APIs to streamline the orchestration process. The application requires a platform that automatically manages underlying infrastructure and scales to zero when there is no activity.
Which primary compute resource should your developers select, given these requirements?

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

Answer: B

Explanation:

Reference: https://firebase.google.com/docs/functions/use-cases
Graphical user interface, text, application, email Description automatically generated

NEW QUESTION 19
- (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 24
- (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 maintence and capacity management cloud provider is the responsible part.

NEW QUESTION 25
- (Topic 1)
Your organization is moving an application to Google Cloud. As part of that effort, it needs to migrate the application's working database from another cloud provider to Cloud SQL. The database runs on the MySQL engine. The migration must cause minimal disruption to users. Data must be secured while in transit.

Which should your organization use?

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

Answer: C

Explanation:

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

NEW QUESTION 28
- (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 30
- (Topic 1)
As your organization increases its release velocity, the VM-based application upgrades take a long time to perform rolling updates due to OS boot times. You need to make the application deployments faster.
What should your organization do?

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

Answer: B

NEW QUESTION 35
- (Topic 1)
An organization wants to dynamically adjust its application to serve different user needs. What are the benefits of storing their data in the cloud for this use case?

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

Answer: C

Explanation:

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

NEW QUESTION 37
- (Topic 1)
Your customer is making a decision on whether to move to Google Cloud. Their key concern is about 10,000 VMs that are part of their IT infrastructure used across more than 110 applications. They are apprehensive of too many changes at this stage. They want to get to Google Cloud in the easiest way possible with minimal disruption. What option would you recommend for them?

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

Answer: D

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

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

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

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

NEW QUESTION 42
- (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 43
- (Topic 1)
A retail store has discovered a cost-effective solution for creating self-service kiosks. They can use existing check-out hardware and purchase a virtual customer service application. Why do they also need an API?

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

Answer: B

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

NEW QUESTION 47
- (Topic 1)
Your organization stores highly sensitive data on-premises that cannot be sent over the public internet. The data must be processed both on-premises and in the cloud.
What should your organization do?

A. Configure Identity-Aware Proxy (IAP) in your Google Cloud VPC network
B. Create a Cloud VPN tunnel between Google Cloud and your data center
C. Order a Partner Interconnect connection with your network provider
D. Enable Private Google Access in your Google Cloud VPC network

Answer: C

Explanation:
After the service provider provisions your connection, you can start passing traffic between your networks by using the service provider's network.
Reference: https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview

NEW QUESTION 52
- (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 57
- (Topic 1)
What would provide near-unlimited availability of computing resources without requiring your organization to procure and provision new equipment?

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

Answer: A

Explanation:

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

NEW QUESTION 60
- (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 65
- (Topic 1)
Your company has been using a shared facility for data storage and will be migrating to Google Cloud. One of the internal applications uses Linux custom images that need to be migrated.
Which Google Cloud product should you use to maintain the custom images?

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

Answer: B

Explanation:

Reference: https://cloud.google.com/compute/docs/images/create-delete-deprecate- private-images
A custom image is a boot disk image that you own and control access to. Use custom images for the following tasks:
Import a virtual disk to Compute Engine from your on-premises environment or from VMs that are running on your local workstation or on another cloud platform.
You can manually import boot disk images to Compute Engine, but one disk at a time.
Graphical user interface, text, application, email Description automatically generated
https://cloud.google.com/compute/docs/images

NEW QUESTION 70
- (Topic 1)
The government has ordered an audit of your company's data. You have hired an external company to conduct the audit. They need to be able to review the data stored in your Cloud Storage buckets across eight projects. How would you grant them access?

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

Answer: B

Explanation:

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

NEW QUESTION 74
- (Topic 1)
Your company security team manages access control to production systems using an LDAP directory group.
How is this access control managed in the Google Cloud production project?

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

Answer: C

Explanation:
Reference: https://cloud.google.com/blog/products/identity-security/achieving-identity-and-access-governance-on-google-cloud

NEW QUESTION 75
- (Topic 1)
An organization wants to move from a strategic cloud adoption maturity level to a trans- formational one. How should the organization change the way they scale?

A. None of these
B. Deploy changes when problems arise.
C. Deploy changes programmatically.
D. Review changes manually.

Answer: C

Explanation:
Because automation is a transformational approach which ensures changes are constant and low-risk.

NEW QUESTION 77
- (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 78
- (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 81
- (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 83
- (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 85
- (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 86
- (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 91
- (Topic 3)
An organization is training a machine learning model to predict extreme weather events in their country.
How should they collect data to maximize prediction accuracy?

A. Collect all weather data evenly across all cities
B. Collect all weather data primarily from at-risk cities
C. Collect extreme weather data evenly across all cities
D. Collect extreme weather data primarily from at-risk cities

Answer: A

Explanation:
 Collect all weather data evenly across all cities. Mainly because it seems that the emphasis for data collection for ML is to make sure there are no holes in your data collection.

NEW QUESTION 96
- (Topic 3)
An organization wants to build autoscaling web applications without having to manage application infrastructure
Which Google Cloud product should they use?

A. App Engine
B. AutoML
C. Anthos
D. Apigee

Answer: A

Explanation:
 Per Google docs, App Engine, allows for "freeing up your developers with zero server management and zero configuration deployments".
https://cloud.google.com/appengine

NEW QUESTION 98
- (Topic 3)

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

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

Answer: C

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

NEW QUESTION 100
- (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 102
- (Topic 3)
An organization wants to use BigQuery data analytics to understand their website performance, but wants to move only some data into the cloud.
Which environment should the organization use?

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

Answer: D

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

NEW QUESTION 104
- (Topic 3)
After rolling out a new update, an organization found a minor bug in its online video game. How should the organization approach this bug while following SRE principles?

A. Accept and learn from the bug because failure is normal
B. Accept and ignore the bug because it is only minor
C. Hold a postmortem to reprimand the employee responsible for the bug
D. Document bug correction to eliminate all future bugs

Answer: A

Explanation:
https://www.blameless.com/sre/sre-principles
Accepting failure as normal is one of the SRE principles. SREs believe that accepting
failure as normal helps to build an iterative, collaborative culture. One way this is done is by holding a blameless "lessons learned" discussion after an incident occurs.

NEW QUESTION 105
- (Topic 3)
How does Google Cloud ensure that customer data remains secure and private when at rest?

A. By aggregating training data for customers within each industry
B. By automatically locking files containing suspicious code
C. By auditing platform privacy practices against industry standards
D. By providing privacy reviews for critical customer applications

Answer: C

Explanation:
Google Cloud commitment to keep the data secure and private:
* 1. Org owns the data and not Google
* 2. Google does not sell data to 3rd parties
* 3. All customer data is encrypted by default
* 4. Google Cloud guards insider against your data
* 5. No backdoor access to any govt. entity
* 6. Google's privacy practices are audited against international standards

NEW QUESTION 109
- (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 113
- (Topic 3)
An organization delivers a proactive healthcare service. They want to efficiently and automatically collect patient data.
What should the organization encourage the patients to do?

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

Answer: B

NEW QUESTION 117
- (Topic 3)
An organization's public cloud provider failed to meet their SLA of 99.99% availability.
What is the potential impact on the organization?

A. The organization risks using up their error budget.
B. Renegotiation of the SLA to put less emphasis on uptime could be necessary.
C. Unexpected downtime could risk the loss of customers.
D. All data stored in their database could be unexpectedly lost.

Answer: C

NEW QUESTION 120
- (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 123
- (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 125
- (Topic 3)
An organization operates their entire IT infrastructure from Google Cloud. What should they do to prepare for data breaches?

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

Answer: C

NEW QUESTION 127

- (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 130
- (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 133
- (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 138
- (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 141
- (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 146
- (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 150
- (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 151
- (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 153
- (Topic 3)
An organization wants to transform multiple types of structured and unstructured data in the cloud from various sources. The data must be readily accessible for analysis and insights.
Which cloud data storage system should the organization use?

A. Relational database
B. Private data center
C. Data field
D. Data warehouse

Answer: D

Explanation:
It supports real-time insights. A data warehouse is an enterprise system used for the analysis and reporting of structured and semi-structured data from multiple sources, https://cloud.google.com/learn/what-is-a-data-warehouse

NEW QUESTION 157
- (Topic 3)
A food delivery service needs access to real-time menu information from all partner restaurants. They also need to share customer order information with the restaurants in real time.
What should the organization use?

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

Answer: B

NEW QUESTION 158
- (Topic 3)
An organization is migrating their business applications from on-premises to the cloud. How could this impact their operations and personnel costs?

A. Reduced on-premises infrastructure management costs
B. Increased on-premises hardware maintenance costs
C. Reduced cloud software licensing costs
D. Increased cloud hardware management costs

Answer: A

NEW QUESTION 160
- (Topic 3)
When an organization adopts cloud technology, how does their total cost of ownership (TCO) shift?

A. Away from cost management toward capital expenditure
B. Away from operational expenditure toward cost management
C. Away from capital expenditure toward operational expenditure
D. Away from operational expenditure toward capital expenditure

Answer: C

NEW QUESTION 161

- (Topic 3)
An organization has created an application that can diagnose different medical conditions when users submit images of their affected body parts.
Which Google Cloud product or service did the organization use?

A. App Engine
B. Machine learning
C. Cloud Logging
D. Cloud Profiler

Answer: B

NEW QUESTION 164
- (Topic 3)
An organization needs to store structured, semi-structured, and unstructured data in its raw, native format in the same repository.
Which cloud data management solution should the organization use?

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

Answer: B

Explanation:
A data lake can store all types of data with no fixed limitation on account size or file and with no specific purpose defined yet. The data comes from disparate sources and can be structured, semi-structured, or even unstructured. Data-lake data can be queried as needed.
https://cloud.google.com/learn/what-is-a-data-lake
A data lake is a centralized repository designed to store, process, and secure large amounts of structured, semistructured, and unstructured data. It can store data in its native
format and process any variety of it, ignoring size limits.

NEW QUESTION 167
- (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 171
- (Topic 3)
An organization is searching for an open-source machine learning platform to build and deploy their own custom machine learning applications using TPUs.
Which Google Cloud product or service should the organization use?

A. TensorFlow
B. BigQuery ML
C. Vision API
D. AutoML Vision

Answer: A

Explanation:
https://en.wikipedia.org/wiki/TensorFlow TensorFlow is a free and open-source software library for machine learning and artificial intelligence. Developer Google Brain Team

NEW QUESTION 173
- (Topic 3)
Several departments in an organization are working together on a project. The organization wants to customize access to resources for each department.
What is the quickest and most efficient way to achieve this?

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

Answer: A

NEW QUESTION 176
- (Topic 3)
An organization wants its users to validate a series of new features for their app. Why should they use App Engine?

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

Answer: C

NEW QUESTION 178
- (Topic 3)
An organization wants to move from a tactical cloud adoption approach to a transformational approach.
How should they adapt the way they lead the organization?

A. Increase top-down visibility and foster a culture of blamelessness
B. Shift from an operational expenditure model to capital expenditure
C. Drive cloud adoption with an individual contributor focus
D. Invest in on-premises infrastructure to redesign relationships between IT and employees

Answer: A

NEW QUESTION 183
- (Topic 3)
An organization is using machine learning to make predictions. One of their datasets
mistakenly includes mislabeled data. How will the prediction be impacted?

A. Increased risk of privacy leaks
B. Increased risk of inaccuracy
C. Decreased model compatibility
D. Decreased model training time

Answer: B

NEW QUESTION 186
- (Topic 3)
What is artificial intelligence?

A. Any system that ingests data in real time
B. Any system that automatically structures data
C. Any system capable of a task that requires smart analytics to generate predictions
D. Any system capable of a task that normally requires human cognition

Answer: D

NEW QUESTION 191
- (Topic 3)
What does Cloud Logging help an organization do?

A. Analyze live source code and log code updates.
B. Deploy infrastructure as code.
C. Analyze logs and accelerate application troubleshooting.
D. Manage storage of custom VM images.

Answer: C

NEW QUESTION 194
- (Topic 3)
Why should an organization consider the total cost of ownership (TCO) when moving from on-premises to the cloud?

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

Answer: C

NEW QUESTION 198
- (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 203

- (Topic 3)
An organization wants to create a new application in the cloud to replace an existing on- premises application.
Which application modernization approach should the organization use?

A. Move the application to the cloud, and then change it.
B. Change their application, and then move it to the cloud.
C. Invent in greenfield.
D. Invent in brownfield.

Answer: D

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

NEW QUESTION 205
- (Topic 3)
An organization meets their service level objective (SLO) of 99 999% ("five nines") How much downtime do their end users experience per year?

A. 5 minutes
B. 500 minutes
C. 5 hours
D. 5 days

Answer: A

NEW QUESTION 209
- (Topic 3)
An organization cannot afford to modernize their infrastructure but they want to process data from their legacy system in a modern platform hosted by a business partner
What solution should the organization choose to make their data accessible?

A. Compute Engine
B. Anthos
C. An application programming interlace
D. Google Kubernetes Engine

Answer: C

NEW QUESTION 210
- (Topic 2)
The customer has applications that do data processing on-premise. They have been built using Ha-doop and Spark. What product should I use on Google Cloud?

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

Answer: A

Explanation:
 Because Dataproc is used to run Hadoop/Spark workloads

NEW QUESTION 213
- (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 215
- (Topic 2)
You are working with a government agency. A web application serves users of the country. It al-lows citizens to receive certain services in providing their national identity. Citizens have com-plained that they are seeing delays in web page loading compared to before. On investigating, they are seeing a lot of spurious traffic coming in from a few IPs which they have identified as for-eign. What should they do?

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

Answer: B

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

NEW QUESTION 220
- (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 224
- (Topic 2)
Considering Different Storage and database options e.g. Cloud Datastore, Cloud SQL, Cloud Stor-age, etc. Which of the following statements is/are correct? (Select two answer)

A. Cloud DataStore and Cloud SQL have Terabytes + and Terabytes Capacity respec-tively.
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 226
- (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 227
- (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 bvary in nature.

NEW QUESTION 230
- (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 233
- (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 234
- (Topic 2)
A financial services company is running an experimental application workload that has a very large number of mathematical calculations involving floating-point numbers. The current application that is running on compute engine is not providing enough speed and throughput. What are the options to increase the processing performance?

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

Answer: D

Explanation:
Compute Engine provides graphics processing units (GPUs) that you can add to your virtual machines (VMs). You can use these GPUs to accelerate specific workloads on your VMs such as machine learning and data processing. https://cloud.google.com/compute/docs/gpus

NEW QUESTION 235
- (Topic 2)

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

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

Answer: C

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

NEW QUESTION 239
- (Topic 2)
Which of the following statements is / are correct about Machine Learning?

A. Machine learning examples include chatbots and automated virtual assistants to automate routine customer service tasks and speed up issue resolution.
B. Machine learning automates the job of building statistical models with Human In-tervention.
C. Robotic process automation (RPA) can not be attached with ML.
D. None of the Above.

Answer: A

Explanation:
Customer service
Machine learning examples include chatbots and automated virtual assistants to automate routine customer service tasks and speed up issue resolution.

NEW QUESTION 244
- (Topic 2)
Which of the followings are core components of Anthos?

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

Answer: D

Explanation:

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

NEW QUESTION 246
- (Topic 2)
A customer has contacted you about migrating to Google Cloud. The customer would like to mi-grate 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 247
- (Topic 2)
When you update the function in firebase by deploying updated code, instances for older versions are cleaned up along with build artifacts in and replaced by new instances.

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

Answer: B

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

NEW QUESTION 251
- (Topic 2)
You are working for a hospital that stores its medical images in an on-premises data room and it is provided that the hospitals want to use Cloud Storage for archival storage of these images. You are required to design and implement a solution where the hospital wants an automated process to up-load any new medical images to Cloud Storage. On the basis of this statements which of the follow-ing statement is correct.

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

Answer: B

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

NEW QUESTION 255
- (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 256
- (Topic 2)
Which of the following statements describe the features of a preemptible VM in-stance? (Select Three Answer)

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

Answer: CDE

Explanation:
Instance is alive for no more than 24 hours, Can be pre-empted with a 30 second notice, Discounted Significantly.
Preemptible VM is an instance that you can create and run at a lower cost than normal instances.
However, Compute Engine might stop (pre-empt) these instances if it requires access to those resources for other tasks. Preemptible instances are excess Compute Engine capacity, so their availability varies with usage.
Live at most 24 hours Can be pre-empted with a 30 second notification via API and are Discounted significantly
Reference link- https://cloud.google.com/compute/docs/instances/preemptible

NEW QUESTION 259
- (Topic 2)
Compute Engine provides machine type recommendations to help you optimize the re- source utilization of your virtual machine (VM) instances. What is this capability called?

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

Answer: C

Explanation:
Compute Engine provides machine type recommendations to help you optimize the resource utilization of your virtual machine (VM) instances. These recommendations are generated automatically based on system metrics gathered by the Cloud Monitoring service over the previous 8 days. Use these recommendations to resize your instance's machine type to use the instance's resources more efficiently. This feature is also known as rightsizing recommendations
Reference link- https://cloud.google.com/compute/docs/instances/apply-machine-type- recommendations-for-instances

NEW QUESTION 263
- (Topic 2)
A customer in the European Union region is very clear that their data should not go outside the Eu-ropean Union. Their end users are spread all over the European U. They have to choose a storage option that serves all the users within Asia via web browsers as quickly as possible. Which storage option will work for them?

Multi-regions

Multi-Region Name	Multi-Region Description
ASIA	Data centers in Asia
EU	Data centers within member states ⧉ of the European Union"
US	Data centers in the United States

A. Cloud Storage with a single region that is known to be within the European U
B. Cloud Filestore is connected to virtual machines which are guaranteed to be within the European U

C. Cloud Storage with the multi-region option of European U
D. Cloud Storage with the dual-region option of European U

Answer: C

Explanation:
Multi-region option will use multiple datacenters that are within the European Union. More regions will also help with lower latency since users are spread across the European U.
https://cloud.google.com/storage/docs/locations#considerations

NEW QUESTION 267
- (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 271
- (Topic 2)
While on-premise, an enterprise had multiple teams, each with its own analytics data store. Attempts to converge the storage for centralized, company-wide analysis failed because of speed and scaling issues. What would be the preferred destination architecture on Google Cloud?

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

Answer: C

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

NEW QUESTION 274
- (Topic 2)
How does a least privilege resource access model contribute to cloud security?

A. Google is responsible for determining access to cloud resources.
B. Employees may only access on-premises software with special permission.
C. Only managers and other senior employees have cloud resource access.
D. Employees only have access to the cloud resources necessary for their job.

Answer: D

Explanation:
This is the definition of a least privilege model.
A supporting principle that helps organizations achieve these goals is the principle of least privilege. The principle of least privilege addresses access control and states that an individual should have only the minimum access privileges necessary to perform a specific job or task and nothing more

NEW QUESTION 278
- (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 282
- (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 283
- (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 285
- (Topic 2)
You are consulting for a client who is migrating to Google Cloud. They presently have a matrix or-ganization. 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 plan-ning for the move?

A. On Google Cloud, create a folder corresponding to each tea
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 as-sign 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 288
- (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 con-trol. They want to modify the reliability to 99.999% (five 9s). What do you tell them? (Choose two answer)

A. Yes, that could be possibl
B. If yes, there will be a significantly higher charge be-cause 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 possibl
E. There is hardly any difference to provide another 0.009% availability.
F. Ask them for the reasonable downtime they are willing to absor
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 292
- (Topic 2)
A client is currently running software on their on-premise systems that is bound by a certain
type of license. They are allowed to run the software on virtualized machines. However, they cannot run them on virtualized machines that are shared by two different companies, teams, or projects. What option do they have on Google Cloud?

A. Google Cloud is a public cloud accessed by multiple customers.
B. Allocate a Bare Metal machine.
C. Setup exclusive login to the VM with self-generated security keys.
D. Allocate sole-tenant nodes

Answer: D

Explanation:
Sole-tenancy lets you have exclusive access to a sole-tenant node, which is a physical Compute Engine server that is dedicated to hosting only your project's VMs. Use sole-tenant nodes to keep your VMs physically separated from VMs in other projects, or to group your VMs together on the same host hardware.
https://cloud.google.com/compute/docs/nodes/sole-tenant-nodes

NEW QUESTION 297
- (Topic 2)
In discussions with a prospective customer who wants to move to Google Cloud to make use of the latest, scalable technologies available therein, you learn that there are very strict regulations concern-ing the storage of data. They only have the approval to store it in their current private data cen-ter. What would you advise them?

A. Retain on-premise itself those portions of data and compute which are under regulatio
B. Take advantage of all the other cloud capabilities for remaining work-loads.
C. It is too risky to touch anything in such a scenari
D. It is best to remain entirely on- premise.
E. Regulations are guideline
F. As long as the data remains encrypted, you can move it anywhere.
G. Petition the government for changes to such regulations as all industries are mov-ing to the public clou
H. Then, when the regulations are eased, move to Google Cloud.

Answer: A

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

NEW QUESTION 300
- (Topic 2)
In Google Cloud IAM: if a policy applied at the project level gives you Owner permissions, your access to an individual resource in that project might be restricted to View permission if someone applies a more restrictive policy directly to that resource. What is correct below the options

A. False
B. None of the above.
C. True
D. Not defined by GCP.

Answer: A

Explanation:
Policies are a union of those applied to resources themselves and those inherited from higher levels in the hierarchy. If a parent policy is less restrictive, it overrides a more restrictive policy applied to the resource. If a parent policy is more restrictive, it does not override a less restrictive policy applied to the resource. Therefore, access granted at a higher level in the hierarchy cannot be taken away by policies applied at a lower level in the hierarchy.

NEW QUESTION 301
- (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 record-ing their inputs in another app. The UX testing is done in the last week of the month. Each month the UX testers could be different. How should the IT team manage the users?

A. Since the app is anyways going to be public, create permanent credentials for the UX testers that they can conveniently use each time.
B. It would be a security issue to have users come and g
C. Recommend that the test-ers be permanently hired to plug the vulnerability issue.
D. It would be a security issue to have users come and g
E. Recommend that the test-ers be permanently hired to plug the vulnerability issue.
F. Create a Group with the permissions required to do the test and record their in-put
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 305
- (Topic 2)
A customer of yours has an SLA with their client that a particular service will respond within 4 sec-onds. The end client has reported that it feels slower. Your engineers do a trial at the client site and notice that there seems to be a delay for many of the requests. It's your team's responsibility to iden-tify the issue quickly within the strict timeline for fixes according to the contract, and then fix it. What should you do?

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

Answer: D

Explanation:
Cloud Trace is a built-in tool in the Operations suite to identify issues like latency.
-> Such fixes are unlikely to change core issues like the service itself being architected or written sub-optimally. Though changes like browser, networking, etc. are helpful, it would be the wrong approach to first recommend that the customer upgrade all their hardware and software.
-> Rewriting code and logging information is going to be time consuming. In general though, logging should always be included in code and it can give good insights. But tracing is way more specific and comprehensive for this requirement.
-> In certain cases, we might identify scaling as the issue. But we should first identify the core problem. So, start with tracing. We can also achieve scale in server-ful technologies.
Reference link- https://cloud.google.com/trace

NEW QUESTION 310
- (Topic 2)
A customer has a tens of applications that are dependent on Oracle databases in their on- premise data centers. The customer wants to migrate to Google Cloud. Their long term goal is to move to other cloud native database technologies. What options do they have to initially move their data?

A. Migrate to a Bare Metal server.
B. Migrate to Cloud SQL.
C. Since there is no hosted Oracle solution, leave the Oracle data on-premise while doing analytics on Google Cloud.
D. Containerize Oracle and run it using Cloud Run.

Answer: B

Explanation:
The Bare Metal solution is the recommended approach. You can deploy Oracle capabilities like clustered databases, replication, and all performance features at licensing costs that are similar to on-premise systems

Choose a Google Cloud bare metal migration strategy 🔖

Send feedback

This article describes the three most common options for migrating your bare metal workloads to Google Cloud along with a framework for understanding your workload requirements. It also explains how to choose the bare metal option that's right for your situation. Finally, it provides practical use cases for each migration strategy.

This article is designed for IT managers and staff who want to understand the capabilities of the Google Cloud offerings Migrate for Compute Engine, Bare Metal Solution, and Mainframe Modernization, and how each can facilitate the migration of bare-metal workloads. The article also discusses an IBM offering for working on Google Cloud.

Migrating to bare metal in Google Cloud serves as an important step toward transforming your IT strategy to focus on the cloud. By running your bare metal workloads closer to Google Cloud services, you can take advantage of those services while implementing your application modernization strategy in parallel.

https://cloud.google.com/architecture/migrating-bare-metal-workloads

NEW QUESTION 312
- (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 317
- (Topic 2)
What according to you are NOT the key capabilities of In-App Messaging?

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

Answer: C

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

NEW QUESTION 322
- (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 usag
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

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

NEW QUESTION 324
- (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 325
- (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 bille
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 bille

F. When failover occurs, connections to the primary instance are automaticallytransferred 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 Consol
I. When failover occurs, connections to the primary instance are automati-cally 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 328
- (Topic 2)
Customer Managed Encryption Keys (CMEK) can be used for encrypting data inside Cloud BigTable, which of the following statements is/are correct. (Select two answer)

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

Answer: BC

Explanation:
Customer-managed encryption keys for Cloud BigTable.
By default, all the data at rest in Cloud Bigtable is encrypted using Google's default encryption. Bigtable handles and manages this encryption for you without any additional action on your part.
If you have specific compliance or regulatory requirements related to the keys that protect your data, you can use customer-managed encryption keys (CMEK) for BigTable. Instead of Google managing the encryption keys that protect your data, your BigTable instance is protected using a key that you control and manage in Cloud Key Management Service (Cloud KMS).
Features
Security: CMEK provides the same level of security as Google's default encryption but provides more administrative control.
Data access control: Administrators can rotate, manage access to, and disable or destroy the key used to protect data at rest in BigTable .
Auditability: All actions on your CMEK keys are logged and viewable in Cloud Logging. Comparable performance: BigTable CMEK-protected instances offer comparable performance to BigTable instances that use Google default encryption.
Flexibility: You can use the same CMEK key in multiple projects or instances or you can use separate keys, depending on your business needs.

NEW QUESTION 329
- (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 330
- (Topic 2)
Your client is a financial services company giving loans based on customer profiles. As part of the regulatory compliance, they have to collect a bunch of different documents with know your cus-tomer (KYC) information. They want to be able to process the information in these documents quickly and at scale. They want to integrate the chosen solution as quickly as possible. What are your options on Google Cloud?

A. Integrate the Cloud Vision API to create a custom model to handle the documents.
B. Create a model using TensorFlow and integrated it into the process workflow.
C. Integrate the Lending DocAI and Document AI in two there processes workflow of the processing loan requests.
D. Integrate the Natural Language API to read the request sent in by clients and to process the forms.

Answer: C

Explanation:
Lending DocAI is a pre-packaged AI solution that speeds "up the mortgage workflow processes to easily process loans and automate document data capture, while ensuring the accuracy and breadth of different documents (e.g., tax statements and asset documents)."
https://cloud.google.com/solutions/lending-doc-ai

NEW QUESTION 334
- (Topic 2)
You are a database manager working for a new product that will need millions of reading and writ-ing 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 follow-ing 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 335
- (Topic 2)
A bank wants to track the success of their existing ATM network, which has been modernized with APIs to instantly notify customers about their transfers. What is the benefit of using Apigee to achieve this goal?

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

Answer: A

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

NEW QUESTION 339
- (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 342
- (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 346
- (Topic 2)
You are working in a company that provides different services to its customer. Now it also wants to offer some paid API services to its B2B customers for e.g. google provides google maps API, cloud vision API, and language translation API. You need to figure out the best solution for the service.

A. Java Programming Spring Boot Framework for to solve the problem of APIs man-agement.
B. Cloud Functions with Firestore and payment gateways integration development.
C. Apigee API Management
D. Frontend & Backend Development with NodeJs and angular etc.

Answer: C

Explanation:
 A top-level idea about Apigee API Management and its offered features can help you solve all questions related to Apigee in Cloud Digital Leader Practice Exam. Apigee is a platform for developing and managing APIs. By fronting services with a proxy
layer, Apigee provides an abstraction or facade for your backend service APIs and provides security, rate limiting, quotas, analytics, and more.
Apigee services: The APIs that you use to create, manage, and deploy your API proxies. Apigee runtime: A set of containerized runtime services in a Kubernetes cluster that Google maintains. All API traffic passes through and is processed by these services.

NEW QUESTION 350
- (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 are inherent to the public cloud provider

NEW QUESTION 354
- (Topic 1)
The CFO is attending one of the preliminary meetings in the migration strategy meeting. She brings up the concern about costs. They have contracts with their vendors and the payments they will need to make when purchasing any kind of infrastructure. This gives them a clear view of numbers for resource budgeting and planning. Can she get the same kind of clarity on Google Cloud?

A. Ye
B. Do a trial run of typical workload
C. See the billing amount and that becomes the base reference.

D. Yes, the Cloud Native Computing Foundation publishes yearly numbers on the cost of running the clou
E. Use that as a reference.
F. Yes, the Pricing Calculator can be used to estimate the cost of resources.
G. Yes, Google provides a typical cost of application workloads by region and indus-tr
H. Use that as a reference.

Answer: C

Explanation:
 The pricing calculator can be used to give clear estimates of resource usage.
-> Running test loads is as closely indicative as using the pricing calculator.
-> There are no cloud cost references published, either by Google or CNCF. Even if some companies have published such info. It might not apply to you.
Reference link:- https://cloud.google.com/products/calculator

NEW QUESTION 359
- (Topic 1)
Your organization needs to categorize objects in a large group of static images using machine learning. Which Google Cloud product or service should your organization use?

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

Answer: C

Explanation:

Reference: https://cloud.google.com/vision
Derive insights from your images in the cloud or at the edge with AutoML Vision or use pre- trained Vision API models to detect emotion, understand text, and more.
Vision API offers powerful pre-trained machine learning models through REST and RPC APIs. Assign labels to images and quickly classify them into millions of predefined categories. Detect objects and faces, read printed and handwritten text, and build valuable metadata into your image catalog.

NEW QUESTION 363
- (Topic 1)
An organization currently stores its data on-premises and they receive different levels of traffic on their website every month. How could the organization benefit from modernizing their infrastructure with cloud technology?

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

Answer: B

Explanation:

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

NEW QUESTION 368
- (Topic 1)
Your company's development team is building an application that will be deployed on Cloud Run. You are designing a CI/CD pipeline so that any new version of the application can be deployed in the fewest number of steps possible using the CI/CD pipeline you are designing. You need to select a storage location for the images of the application after the CI part of your pipeline has built them.
What should you do?

A. Create a Compute Engine image containing the application
B. Store the images in Container Registry
C. Store the images in Cloud Storage
D. Create a Compute Engine disk containing the application

Answer: B

Explanation:

Reference: https://cloud.google.com/container-registry/docs/pushing-and-pulling

NEW QUESTION 369
- (Topic 1)
An organization has created an ecommerce website. What data on this website would be considered structured data?

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

Answer: D

Explanation:
Because product ratings are structured because they are numerical scores.

NEW QUESTION 373
- (Topic 1)
Your organization consists of many teams. Each team has many Google Cloud projects. Your organization wants to simplify the management of identity and access policies for these projects.
How can you group these projects to meet this goal?

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

Answer: C

Explanation:

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

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

NEW QUESTION 374
- (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 377
- (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 380
- (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 383
- (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 387
- (Topic 1)
Which Google Cloud product gives you a consistent platform for multi-cloud application deployments and extends other Google Cloud services to your environment?

A. Google Kubernetes Engine
B. Virtual Public Cloud
C. Compute Engine
D. Anthos

Answer: D

Explanation:

https://cloud.google.com/anthos

NEW QUESTION 388
- (Topic 1)
Which Google Cloud product or feature makes specific recommendations based on security risks and compliance violations?

A. Google Cloud firewalls

B. Security Command Center
C. Cloud Deployment Manager
D. Google Cloud Armor

Answer: B

Explanation:

Reference: https://cloud.google.com/security-command-center
Security Command Center is Security and risk management platform for Google Cloud.
Reference link- https://cloud.google.com/security-command-center

NEW QUESTION 390
- (Topic 1)
A customer has new applications to build that has to handle both batch data and streaming data. Which product should they choose?

A. Dataprep
B. Dataflow
C. Dataproc
D. Data Fusion

Answer: B

Explanation:
 Enabling Requester Pays is useful, for example, if you have a lot of data you want to make available to users, but you don't want to be charged for their access to that data.
Reference link- https://cloud.google.com/storage/docs/requester-pays

NEW QUESTION 393
- (Topic 1)
Your organization is developing an application that will manage payments and online bank accounts located around the world. The most critical requirement for your database is that each transaction is handled consistently. Your organization anticipates almost unlimited growth in the amount of data stored.
Which Google Cloud product should your organization choose?

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

Answer: D

Explanation:
Features of Cloud Spanner
Reference: https://k21academy.com/google-cloud/cloud-sql-vs-cloud-spanner/

NEW QUESTION 395
- (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 400
......