ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 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

Question: 2

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

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

Answer: C

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

Question: 3

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

Question: 4

The operating systems of some of your organization's virtual machines may have a security vulnerability.

How can your organization most effectively identify all virtual machines that do not have the latest security update?

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

Answer: A

Explanation:

Security Health Analytics and Web Security Scanner detectors generate vulnerabilities findings that are available in Security Command Center. Your ability to view and edit findings is determined by the

Identity and Access Management (IAM) roles and permissions you are assigned. For more information about IAM roles in Security Command Center.

Reference link:-

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

Question: 5

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

What should you do?

A. Renew your licenses for an additional period of 3 years. Renew your licenses for an additional period of 3 years. Negotiate a cost reduction with your current hosting provider wherein infrastructure cost is reduced when workloads are not in use

B. Renew your licenses for an additional period of 2 years. Negotiate a cost reduction by committing to an automatic renewal of the licenses at the end of the 2 year period
C. Migrate the workloads to Compute Engine with a bring-your-own-license (BYOL) model
D. Migrate the workloads to Compute Engine with a pay-as-you-go (PAYG) model

Answer: D

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

Question: 6

Your organization runs a distributed application in the Compute Engine virtual machines. Your organization needs redundancy, but it also needs extremely fast communication (less than 10 milliseconds) between the parts of the application in different virtual machines.
Where should your organization locate this virtual machines?

A. In a single zone within a single region
B. In different zones within a single region
C. In multiple regions, using one zone per region
D. In multiple regions, using multiple zones per region

Answer: B

Explanation:
Multi zone is also redundant within the region and it provides the lowest latency.

Reference link:-
https://cloud.google.com/solutions/best-practices-compute-engine-region-selection

Question: 7

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: A, D

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

Question: 8

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

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

Answer: A

Explanation:

What is a preemptible instance?

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

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

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

Question: 9

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

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

Answer: A

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

Question: 10

Your multinational organization has servers running mission-critical workloads on its premises around the world. You want to be able to manage these workloads consistently and centrally, and you want to stop managing infrastructure.
What should your organization do?

A. Migrate the workloads to a public cloud
B. Migrate the workloads to a central office building
C. Migrate the workloads to multiple local co-location facilities
D. Migrate the workloads to multiple local private clouds

Answer: A

Explanation:
Only public cloud offers to centrally manage the infra. for Pvt cloud it may not be possible to get same Pvt Cloud provider across the globe.

Question: 11

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

Question: 12

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

Question: 13

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

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

Answer: B

Explanation:

Benefits of SaaS

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

What are laas, Paas and Saas?

laaS. PaaS and SaaS are the three most popular types of cioud service offerings. (They are sometimes referred to as cloud service models or cloud computing service models.)

- laaS. or nfrastructu-e as a service, s on-demand access to cloud-hosted physical and virtual servers, storage and network ng - the tackend IT infrastructure for running applications and workloads In the cloud

- PaaS. or platform as a service, Is on-demand access to a complete, reao'y-to- use. cloud-hosted platform for developing, tunning, maintaining and maraging applications.

- SaaS, or software as a serv ce. is on-demand access to ready-to-use, cloud- hosted application software.

laaS. PaaSand SaaS are not mutually exclusive. Many mid-sized bus nessesuse more than one, and most large enterprises use all th rec

Asaservice' refers to the way IT assets are consumed in these offer ngs rand-tothe essential d fference between cloud comoutng and traditional IT. In traditional IT. ar organization consumes IT assets - hardware, system software, development tools, applications - by purchasing them, installing them, managing them and maTta nirg them i r its own on-premises data center. In cioud com puting.

the cloud service provider owns, manages and maintains the assets; the customer consumes them via an Internet connection, and pays for them on a subscription or pay-as-you-go basis.

Question: 14

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

Explanation:

Question: 15

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

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

Answer: D

Explanation:

Configure single sign-on in the Google domain

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

Federating Google Cloud with Active Directory □

Send feedback

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

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

This article assumes that you're familiar with Active Directory

Implementing federation

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

https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-introduction

Reference Link- https://cloud.google.com/architecture/identity/single-sign-on

Question: 16

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

Question: 17

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/

Google Cloud SQL is a fully managed service offered by Google Cloud Platform Google Cloud SQL is a MySQL database inside Google Cloud. There is no need to install, maintain and create admin accounts because it is fully managed by Google Cloud. It helps you create, modify, configure and utilize a relational database, same as MySQL.

Google sends constant updates and adds new features to its services to fulfil the business requirements of its users

Let's have a look at the improvements done in Google Cloud SQL.

- Google deliberately increased the storage space to 100 Gigabytes which was 10 Gigabytes earlier.
- The modified version is loaded with the capacity of 16GB RAM to run instances hassle-free.
- Increased RAM has helped users to keep four times more cache than earlier.
- Now Google provides both replicated and non-replicated databases

Question: 18

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:

E. 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, ant 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.

Question: 19

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

BigQuery ML and AI Platform □

Learn how to build a system to predict customer propensity to purchase by using BigQuery ML and AI Platform.

You can use a propensity to purchase system to predict customers who are most likely to make a purchase, so that you can personalize communications with them. Use online predictions to take real-time action based on user behavior on your website, or batch predictions to inform less time-sensitive communications like email.

This tutorial shows you how to create a logistic regression □ model to determine whether a customer will make a purchase. This type of model is used because it is good for evaluating the probability of an outcome. The model evaluates metrics that reflect customer behavior on a website, and assigns the customer a probability to purchase value between 0 and 1 based on this data. The model then sets a label indicating "likely to purchase" for any customer with a probability of greater than .5.

This tutorial uses the Google Analytics Sample and ecommerce datasets to train the model. These datasets are hosted publicly on BigQuery. These datasets provide 12 months (August 2016 to August 2017) of obfuscated Analytics 360 data from the Google Merchandise Store, a real e-commerce store that sells Google-branded merchandise.

To apply the lessons from this tutorial to a production use case, you could use your own Analytics 360 data, or data from a similar system that gives you access to metrics about customer behaviour on your website.

Question: 20

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

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

Answer: D

Explanation:
Reference: https://cloud.google.com/bigtable
Cloud Bigtable is a sparsely populated table that can scale to billions of rows and thousands of columns, enabling you to store terabytes or even petabytes of data. A single value in each row is indexed; this value is known as the row key. Bigtable is ideal for storing very large amounts of singlekeyed data with very low latency. It supports high read and write throughput at low latency, and it is an ideal data source for MapReduce operations.
Bigtable is exposed to applications through multiple client libraries, including a supported extension to the Apache HBase library for Java. As a result, it integrates with the existing Apache ecosystem of open-source Big Data software.
Bigtable's powerful back-end servers offer several key advantages over a self-managed HBase installation:

Incredible scalability. Bigtable scales in direct proportion to the number of machines in your cluster. A self-managed HBase installation has a design bottleneck that limits the performance after a certain threshold is reached. Bigtable does not have this bottleneck, so you can scale your cluster up to handle more reads and writes.
Simple administration. Bigtable handles upgrades and restarts transparently, and it automatically maintains high data durability. To replicate your data, simply add a second cluster to your instance, and replication starts automatically. No more managing replicas or regions; just design your table schemas, and Bigtable will handle the rest for you.
Cluster resizing without downtime. You can increase the size of a Bigtable cluster for a few hours to handle a large load, then reduce the cluster's size again—all without any downtime. After you change a cluster's size, it typically takes just a few minutes under load for Bigtable to balance performance across all of the nodes in your cluster.

Question: 21

Your organization needs to build streaming data pipelines. You don't want to manage the individual servers that do the data processing in the pipelines. Instead, you want a managed service that will automatically scale with the amount of data to be processed.
Which Google Cloud product or feature should your organization choose?

A. Pub/Sub
B. Dataflow

C. Data Catalog
D. Dataprep by Trifacta

Answer: B

Explanation:
Reference: https://cloud.google.com/dataflow/docs/guides/deploying-a-pipeline

Deploying a pipeline n

Send feedback

★ This document explains in detail how Dataflow deploys and runs a pipeline, and covers advanced topics like optimization and load balancing. If you are looking for a step-by-step guide on how to create and deploy your first pipeline, use Dataflow's quickstarts for Java, Python, Go, or templates

After you construct and test your Apache Beam pipeline, you can use the Dataflow managed service to deploy and execute it Once on the Dataflow service, your pipeline code becomes a Dataflow job

The Dataflow service fully manages Google Cloud services such as Compute Engine and Cloud Storage to run your Dataflow job. automatically spinning up and tearing down the necessary resources. The Dataflow service provides visibility into your job through tools like the Dataflow monitoring interface and the Dataflow command-line interface

★ You can control some aspects of how the Dataflow service runs your job by setting execution parameters in your pipeline code. For example, the execution parameters specify whether the steps of your pipeline run on worker virtual machines, on the Dataflow service backend, or locally.

In addition to managing Google Cloud resources, the Dataflow service automatically performs and optimizes many aspects of distributed parallel processing. These include the following

- Parallelization and distribution. Dataflow automatically partitions your data and distributes your worker code to Compute Engine instances for parallel processing. For more information, see parallelization and distnbution.
- Optimization. Dataflow uses your pipeline code to create an execution graph that represents your pipeline's PCollection s and transforms, and optimizes the graph for the most efficient performance and resource usage. Dataflow also automatically optimizes potentially costly operations, such as data aggregations. For more information, see Fusion optimization and Combine optimization

Reference link- https://cloud.google.com/dataflow/docs/guides/deploying-a-pipeline

Question: 22

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

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

Answer: C

Explanation:

Deploy your application to App Engine using the gcloud app deploy command. This command automatically builds a container image by using the Cloud Build service and then deploys that image to the App Engine flexible environment.
Reference: https://cloud.google.com/appengine/docs/flexible/nodejs/testing-and-deploying-your- app

Question: 23

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

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

Answer: D

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

1. Gain centralized visibility and control
2. Discover misconfigurations and vulnerabilities
3. Report on and maintain compliance
4. Detect threats targeting your Google Cloud assets

https://cloud.google.com/security-command-center

Question: 24

Your organization needs to establish private network connectivity between its on-premises network and its workloads running in Google Cloud. You need to be able to set up the connection as soon as possible.
Which Google Cloud product or feature should you use?

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

Answer: A

Explanation:

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

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

Question: 25

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

Question: 26

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

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

Question: 27

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

About Google Cloud Directory Sync

With Google Cloud Directory Sync (GCDS), you can synchronize the data in your Google Account with your Microsoft Active Directory or LDAP server. GCDS doesn't migrate any content (such as email messages, calendar events, or files) to your Google Account. You use GCDS to synchronize your Google users, groups, and shared contacts to match the information in your LDAP server.

https://support.google.com/a/answer/106368?hl=en

Question: 28

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

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

Answer: C

Explanation:

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

Question: 29

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:

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

Question: 30

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

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

Answer: D

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

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

Question: 31

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/

Question: 32

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

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

Answer: D

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

Question: 33

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

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

Answer: D

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

Question: 34

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

Question: 35

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 enhanced support ensuring the long term success of your SecureAuth SaaS IAM deployment

SecureAuth is dedicated to providing the **Industry-leading** enhanced **support** ensuring the long term success of your SecureAuth SaaS 1AM deployment

While our basic support offers industry leading coverage and response times for some customers, SecureAuth protects critical applications meaning delays and extended downtime is simply not an option. For these customers our **Enhanced Support** offerings provide **24*7 coverage** and the most responsive and **complete** SLAs available. That's why we offer three different levels of support, so you can choose the level of support that best works for your needs.

Question: 36

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

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

Answer: B

Explanation:

An identity platform is a modern solution for managing the identities of users and devices in a centralized fashion.
Reference: https://www.okta.com/blog/2021/07/what-is-an-identity-

platform/#:~:text=An%20identity%20platform%20is%20a,%2C%20integrations%2C%20and%20platfo
rm%20services

An identity platform is a modern solution for managing the identities of users and devices in a centralized fashion. It enables organizations to securely authorize workforce and customer users to access their ecosystem using access management tools, programmable components, integrations, and platform services.

Today's organizations have a wide range of identity requirements. As they expand, embrace new innovations, and meet new customer demands, they need an identity solution that can grow with them. And as they grow, these organizations are also looking for offerings that centralize and consolidate identity, reducing the need for multiple access management, governance, and authentication products that don't necessarily talk to each other.

In short, although the concept of identity platforms is still new, the growth of disruptive technologies and highly personalized products and services has made a platform approach the key for effectively unifying identity management. But how did we get here?

Question: 37

Which Google Cloud service or feature lets you build machine learning models using Standard SQL and data in a data warehouse?

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

Answer: A

Explanation:

BigQuery ML lets you create and execute machine learning models in BigQuery using standard SQL queries.
Reference: https://cloud.google.com/bigquery-
ml/docs/introduction#:~:text=BigQuery%20ML%20lets%20you%20create,the%20need%20to%20mo ve%20data

https://cloud.google.com/bigquery-ml/docs/introduction

Question: 38

Your organization runs an application on virtual machines in Google Cloud. This application processes incoming images. This activity takes hours to create a result for each image. The workload for this application normally stays at a certain baseline level, but at regular intervals it spikes to a much greater workload. Your organization needs to control the cost to run this application.
What should your organization do?

A. Purchase committed use discounts for the baseline load
B. Purchase committed use discounts for the expected spike load
C. Leverage sustained use discounts for your virtual machines
D. Run the workload on preemptible VM instances

Answer: C

Explanation:

The idea of the Sustained Use discount is that the longer you run a VM instance in any given month, the bigger discount you will get from the list price.
Reference: https://www.parkmycloud.com/blog/google-sustained-use-discounts/

Question: 39

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

Question: 40

Your organization runs many workloads in different Google Cloud projects, each linked to the same billing account. Each project's workload costs can vary from month to month, but the overall combined cost of all projects is relatively stable. Your organization needs to optimize its cost. What should your organization do?

A. Purchase a commitment per project for each project's usual minimum
B. Create a billing account per project, and link each project to a different billing account
C. Turn on committed use discount sharing, and create a commitment for the combined usage
D. Move all workloads from all different projects into one single consolidated project

Answer: C

Explanation:

Turn on committed use discount sharing, and create a commitment for the combined usage

Sharing your committed use discounts across all your projects reduces the overhead of managing discounts on a per-project basis, and maximizes your savings by pooling all your discounts across your projects' resource usage. If you have multiple projects that share the same Cloud Billing account, you can enable committed use discount sharing so all of your projects within that Cloud Billing account share all of your committed use discount contracts. Your sustained use discounts are also pooled at the same time. That is, sustained use discounts are calculated using the total resources across these projects, rather than just the resources within a single project.

Sharing committed use discounts across projects

Sharing your committed use discounts across all your projects reduces the overhead of managing discounts on a perproject basis, and maximizes your savings by pooling all your discounts across your projects resource usage

If you have multiple projects that share the same Cloud Billing account, you can enable committed use discount sharing so all of your projects within that Cloud Billing account share all of your committed use discount contracts Your sustained use discounts are also pooled at the same time. That is, sustained use discounts are calculated using the total resources across these projects, rather than just the resources within a single project.

For example, if you purchase two commitment contracts for a total of 160 cores, and you run 200 cores during the month, you will receive committed use discounts for 160 cores across the projects that used them. The additional 40 cores will be billed at on-demand, non-committed use rates. After you purchase a set amount of commitments, you're billed forthose commitments monthly, even if you don't use them. For example, if you purchase commitments for 160 cores, you're billed the committed use rates forthose 160 cores for the whole month, even if don't use them See Understanding discount sharing for cost-saving utilization recommendations.

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

Question: 41

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.

Question: 42

Your company has recently acquired three growing startups in three different countries. You want to reduce overhead in infrastructure management and keep your costs low without sacrificing security and quality of service to your customers.

How should you meet these requirements?

A. Host all your subsidiaries' services on-premises together with your existing services.
B. Host all your subsidiaries' services together with your existing services on the public cloud.
C. Build a homogenous infrastructure at each subsidiary, and invest in training their engineers.
D. Build a homogenous infrastructure at each subsidiary, and invest in hiring more engineers.

Answer: B

Explanation:
Host all your subsidiaries' services together with your existing services on the public cloud.

Question: 43

What is the difference between Standard and Coldline storage?

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

Answer: D

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

Question: 44

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

Question: 45

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

Cloud Functions gives developers access to Firebase and Google Cloud events, along with scalable computing power to run code in response to those events. While it's expected that Firebase apps will use Cloud Functions in unique ways to meet their unique requirements, typical use cases might fall into these areas:

- Notify users when something interesting happens.
- Perform database sanitization and maintenance.
- Execute intensive tasks in the cloud instead of in your app.
- Integrate with third-party services and APIs.

Review the use cases and examples for each category that interests you, and then proceed to our Get Started tutorial or to specific how-to guides for authentication events, analytics events, and more. See the eventType API reference for the complete list of supported event types.

Notify users when something interesting happens

Developers can use Cloud Functions to keep users engaged and up to date with relevant information about an app Consider, for example, an app that allows users to follow one another's activities in the app. Each time a user adds themselves as a follower of another user, a write occurs In the Realtime Database Then this write event could trigger a function to create Firebase Cloud Messaging (FCM) notifications to let the appropriate users know that they have gained new followers.

Question: 46

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.

Images □

Use operating system images to create boot disks for your instances. You can use one of the following image types

- Public images are provided and maintained by Google, open source communities, and third-party vendors By default, all Google Cloud projects have access to these images and can use them to create instances.
- Custom images are available only to your Cloud project. You can create a custom image from boot disks and other Images Then, use the custom image to create an instance

You can use most public images at no additional cost, but there are some premium images that do add additional cost to your instances. Custom images that you import to Compute Engine add no cost to your Instances, but do incur an image storage charge while you keep your custom image in your project

Some images are capable of running containers on Compute Engine.

To view the source image for a VM, see Viewing source image.

Public images

Compute Engine offers many preconfigured public images that have compatible Linux or Windows operating systems. Use these operating system images to create and start instances. Compute Engine uses your selected image to create a persistent boot disk for each instance. By default, the boot disk for an instance is the same size as the image that you selected. If your instance requires a larger persistent boot disk than the image size, resize the boot disk.

https://cloud.google.com/compute/docs/images

Question: 47

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

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

Answer: D

Explanation:

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

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

Question: 48

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.

Question: 49

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.

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

Question: 50

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

A. Cloud Storage

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

Answer: C

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

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

Question: 51

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

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

Answer: A

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

Where you're moving data from	Scenario	Suggested products
Another cloud provider (for example, Amazon Web Services or Microsoft Azure) to Google Cloud	-	Storage Transfer Service
Cloud Storage to Cloud Storage (two different buckets)	-	Storage Transfer Service
Your private data center to Google Cloud	Enough bandwidth to meet your project deadline for less than 1 TB of data	gsutil
Your private data center to Google Cloud	Enough bandwidth to meet your project deadline for more than 1 TB of data	Storage Transfer Service for on-premises data
Your private data center to Google Cloud	Not enough bandwidth to meet your project deadline	Transfer Appliance

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

Question: 52

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.

Question: 53

Your organization wants to run a container-based application on Google Cloud. This application is expected to increase in complexity. You have a security need for fine-grained control of traffic between the containers. You also have an operational need to exercise fine-grained control over the application's scaling policies.

What Google Cloud product or feature should your organization use?

A. Google Kubernetes Engine cluster
B. App Engine
C. Cloud Run
D. Compute Engine virtual machines

Answer: A

Explanation:
Google Kubernetes Engine GKE seems a better fit since the requirement is for "security need for finegrained control of traffic between the containers" and "fine-grained control over scaling policies". Such level of control is easier on GKE that Cloud Run.

When it comes to managed Kubernetes services, Google Kubernetes Engine (GKE) is a great choice if you are looking fora container orchestration platform that offers advanced scalability and configuration flexibility. GKE gives you complete control over every aspect of container orchestration, from networking, to storage, to how you set up observability—in addition to supporting stateful application use cases. However, if your application does not need that level of cluster configuration and monitoring, then fully managed Cloud Run might be the right solution for you.

Fully managed Cloud Run is an ideal serverless platform for stateless containerized microservices that don't require

Kubernetes features like namespaces, co-location of containers in pods (sidecars) or node allocation and management.

Reference link- https://cloud.google.com/blog/products/containers-kubernetes/when-to-use- google-kubernetes-engine-vs-cloud-run-for-containers

Question: 54

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

Question: 55

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.

Question: 56

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.

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

Question: 57

Your organization wants to migrate your on-premises environment to Google Cloud. The on-premises environment consists of containers and virtual machine instances. Which Google Cloud products can help to migrate the container images and the virtual machine disks?

A. Compute Engine and Filestore
B. Artifact Registry and Cloud Storage
C. Dataflow and BigQuery
D. Pub/Sub and Cloud Storage

Answer: A

Explanation:
Reference: https://cloud.google.com/compute/docs/import/importing-virtual-disks

If you have virtual disks in your on-premises environment with software and configurations that you need (sometimes referred to as *golden disks* or *golden images*), you can save time by importing those virtual disks Into Compute Engine and using the resulting image to create virtual machines The import tool supports most virtual disk file formats, including VMDK and VHD.

If you exported your disk from Compute Engine, you can create images from the disk

For information about how to create an automated system for migrating several virtual machines (VMs), see Migrating VMs to Compute Engine.

Question: 58

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

When businesses shift from solely on-premises deployments to using cloud-based services, identity management can become more complex. This is especially true when it comes to hybrid and multi-cloud identity management.

Cloud Identity and Access Management (1AM) offers several ways to manage identities and roles in Google Cloud. One particularly important identity management task is identity and access governance (IAG): ensuring that your identity and access permissions are managed effectively, securely, and correctly. A major step in achieving IAG is designing an architecture that suits your business needs and also allows you to satisfy your compliance requirements. To manage the entire enterprise identity lifecycle you must consider the following core tasks:

Question: 59

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

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

Answer: AB

Explanation:

A label is a key-value pair that helps you organize your Google Cloud resources. You can attach a label to each resource, then filter the resources based on their labels. Information about labels is forwarded to the billing system, so you can break down your billed charges by label.

Reference link- https://cloud.google.com/cost-management

Question: 60

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

Question: 61

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:

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

Question: 62

Your organization is defining the resource hierarchy for its new application in Google Cloud. You need separate development and production environments. The production environment will be deployed in Compute Engine in two regions. Which structure should your organization choose?

A. Create a single project for all environments. Use labels to segregate resources by environment.
B. Create a single project for all environments. Use tags to segregate resources by environment.
C. Create one project for the development environment and one project for the production environment.
D. Create two projects for the development environment and two projects for the production environment (one for each region).

Answer: C

Explanation:
Many organizations have separate development and production environments so they can build and test new features without disturbing production traffic. In Optimizely, you can create separate projects for each environment to help with governance.
With separate development and production projects, your organization can safely build and QA experiments and Personalization campaigns in a development environment before deploying to production. This approach allows multiple stakeholders in your organization to act as gatekeepers for running new experiments in production.

Reference link- https://support.optimizely.com/hc/en-us/articles/4410284353805-Set-up-projects- for-development-and-production-environments

Question: 63

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.
D. 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

Question: 64

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

Question: 65

Your team is publishing research results and needs to make large amounts of data available to other researchers within the professional community and the public at minimum cost.

How should you host the data?

A. Use a Cloud Storage bucket and enable "Requester Pays.'
B. Use a Cloud Storage bucket and provide Signed URLs for the data files.
C. Use a Cloud Storage bucket and set up a Cloud Interconnect connection to allow access to the data.
D. Host the data on-premises. and set up a Cloud Interconnect connection to allow access to the data.

Answer: A

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

Question: 66

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

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

Answer: A

Explanation:
Reference: https://cloud.google.com/security/infrastructure/design

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

Question: 67

How do Migrate for Compute Engine and Migrate for Anthos differ?

A. Unlike Migrate for Anthos, Migrate for Compute Engine assumes that the migration source is VMware vSphere.
B. Migrate for Compute Engine charges for ingress, but Migrate for Anthos does not.
C. Migrate for Compute Engine is closed source, and Migrate for Anthos is open source.
D. Migrate for Anthos migrates to containers, and Migrate for Compute Engine migrates to virtual machines.

Answer: D

Explanation:
Reference: https://cloud.google.com/migrate/anthos
Migrate workloads to Compute Engine with Migrate for Compute Engine. Migrate from Compute Engine to containers with Migrate for Anthos and GKE.
This method makes sense, for instance, in cases where you want to conduct a data-center migration and migrate all workloads into Compute Engine, and only at a second stage selectively modernize suitable workloads to containers.

https://cloud.google.com/migrate/containers/docs/architecture

Question: 68

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.

Question: 69

An organization runs their application on a virtual machine, but every time they want to edit specific features, they have to bring the system ofline 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.

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

Question: 70

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

A. Enable committed use with discount sharing for the projects

B. Review the usage of resources by each project on a daily basis. Enable committed use for the following day based on that number, so that it gives maximum granularity without wastage.

C. Take a report of each project's use in the last year. Enable committed use on a per-project basis based on the previous year's numbers.

D. Share a Google Sheet and request each project team to send in their estimate. Enable committed use accordingly on a per-project basis.

Answer: A

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

Question: 71

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

Question: 72

Which of the following NIST Cloud characteristics uses the business model of shared re-sources in a cloud environment?

A. Elasticity
B. Availability
C. Broad Network Access
D. Multi-Tenancy

Answer: D

Explanation:
Explanation

In cloud computing, multitenancy means that multiple customers of a cloud vendor are using the same computing resources. Even though they share resources, cloud customers aren't aware of each other, and their data is kept totally separate. Mul-titenancy is a crucial component of cloud computing; without it, cloud services would be far less practical. Multitenant architecture is a feature in many types of public cloud com-puting, including IaaS, PaaS, SaaS, containers, and serverless computing.

Question: 73

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:

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

Question: 74

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

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

Answer: C

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

Question: 75

A company with its own private data center has called you in for help with their disaster recovery planning. News of multiple ransomware attacks has made them very anxious. They want to make they are well prepared for such an eventuality. Which of these would be good recommendations?

A. It is better to have redundancy; so, set up another private data center nearby so that you can quickly go over in case of an emergency.
B. It is better to have redundancy; use one or many of the Google Cloud datacenters as a backup location.
C. The one data center is enough, as long as the data is encrypted; attackers won't be able to read the data.
D. The one data center is enough as long as you regularly back up data and save it in another place in the same DC.

Answer: B

Explanation:

Explanation
A single data center is vulnerable. So any option involving that is not good.
Reference Link:- https://www.coresite.com/blog/data-center-redundancy

Question: 76

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

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

Answer: C

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

Question: 77

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

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

Answer: C

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

Question: 78

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:

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.

Question: 79

An organization has completely migrated all their infrastructure to the cloud to benefit from its agility. Now they want to innovate faster and achieve a higher return on investment. What should the organization do?

A. Manually provision all cloud infrastructure for increased control.
B. Modernize their applications.
C. Lower their service level objective (SLO).
D. Move to a hybrid architecture with some of their infrastructure on-premises.

Answer: B

Explanation:

Explanation

Because this will enable the business to better serve their users.

Question: 80

Your customer currently has a hybrid cloud setup including their on-premises data center and AWS. They are consolidating all their services on Google Cloud as part of a modernization plan and want to spend less IT effort in the future. There are about 10 MySQL and 25 PostgreSQL databases across the two DCs. What is the best option to for them?

A. Use the Data Catalog Service to manage the metadata of the databases
B. Use Cloud Dataflow service and setup Google's Cloud SQL as the sink and the others as the source, which will cause the data to flow in as expected.
C. Use the Database Migration Service
D. Use the Bare Metal Solution and copy the databases directly as they are on-premises and on AWS.

Answer: C

Explanation:

Explanation
Database Migration is the right one to use: "Simplifying migrations to Cloud SQL. Now available for MySQL and PostgreSQL migrations, with SQL Server coming soon." Since the customer also doesn't want to manage their own database installations in the future, Cloud SQL is the best option.

https://cloud.google.com/database-migration

Question: 81

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:

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.

Question: 82

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.

Question: 83

Your company provides car maintenance services. It is conducting an internal hackathon to identify new ideas that could expand their business. The teams have pitched different ideas and have started working on it. They have to present their application to the judges within 48 hours. A presentation alone is not enough; they have to demonstrate a working proof of concept. The team that you are mentoring is going to recommend additional ser-vices to drive in customers based on the brand of car they drive in. They need to be able to identify what brand of car the customer has, based on a photograph automatically taken at entry. They have already discovered an open source database of car images collected by online enthusiasts. How should they implement this solution?

A. Use Deep Learning Containers that are preconfigured and optimized containers for deep learning environments.
B. Use AutoML Image - upload the images and let it create a working model for you.
C. Use TensorFlow to create a model that will identify the car brands; use the available data to train the model.
D. Use Cloud Vision AI that is able to detect logos. Write only the code to integrate in-to your workflow.

Answer: B

Explanation:

Explanation

It would be most straightforward to use AutoML Image. Put the images in Cloud Storage, point to it from AutoML, and start the model building process.
Reference Link- https://cloud.google.com/automl

Question: 84

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

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

Answer: A

Explanation:

Explanation

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

Question: 85

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

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

Answer: C

Explanation:

Explanation

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

Question: 86

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

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

Answer: B

Explanation:

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

Question: 87

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:
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.

Question: 88

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

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

Answer: D

Explanation:
Explanation
Premium Support will have a 15-minute response time with 24/7 response for high & criti-cal-impact issues.

https://cloud.google.com/support

Question: 89

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.

Question: 90

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

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

Answer: D

Explanation:

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

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

- Security and encryption
- Backups

Software and applications, including:

- Installation
- Configuration
- Upgrades and patching

Operating system and any hypervisor, including:

- Configuration changes
- Upgrades and patching

Server clusters, including:

- Installation
- Configuration
- Maintenance

Licensing

Option B & C is also true.

With Bare Metal Solution, Google Cloud provides and manages the core infrastructure, the net-work, the physical and network security, and hardware monitoring capabilities in an environment from which you can access all of the Google Cloud services. The core infrastructure includes secure, controlled-environment facilities, and power.
The Bare Metal Solution also includes the provisioning and maintenance of the custom, sole-tenancy hardware with local SAN, and smart hands support.
The network, which is managed by Google Cloud includes a low-latency Cloud Interconnect connection into the customer Bare Metal Solution environment.
The available Google Cloud services include private API access, management tools, support, and billing.

Question: 91

An organization has had a data leak scare because one employee made a sensitive Cloud Storage bucket available to the public. Given the nature of the company's business, it is understood that there is never any reason to give the public direct access to any file. The security head wants to ensure that such an event never occurs again. How can you ensure this?

A. Remove Edit access rights of all Cloud Storage buckets so that no user can make any edits.
B. Set an organizational policy constraint to restrict bucket access set to the public.
C. Use Cloud Scheduler to run a job at a specified interval to scan buckets. Any public permissions can be programmatically changed.
D. Write Cloud Functions code connected to Cloud Storage. Any changes will be notified to the function which can be used to reset the public access.

Answer: B

Explanation:
The straightforward way to set it is using Organizational Policy constraint. Any attempts to change the organizational setting will be rejected for any project and resource.

Reference link:

-> https://cloud.google.com/resource-manager/docs/organization-policy/overview
-> https://cloud.google.com/resource-manager/docs/organization-policy/org-policy-constraints

Question: 92

Which of the following options is/are correct about Preemptible VMs?

A. Preemptible VMs don't have fixed pricing.
B. Both A & B
C. None of the Above.
D. You can not use Preemptible VMs at fault-tolerant workloads such as high-performance computing, big data and analytics, continuous integration/continuous delivery (CI/CD), rendering/transcoding, and testing.

Answer: C

Explanation:

Preemptible VMs:

Predictable and low cost
Preemptible VMs are up to 80% cheaper than regular instances. Pricing is fixed so you will always get low cost and financial predictability, without worrying about variable market pricing.

Expand your batch processing
Supplement your regular VMs with lower-cost, preemptible instances to finish your computeintensive work faster, saving you time and money. Throw preemptible VMs at fault-tolerant workloads such as high performance computing, big data and analytics, continuous integration/continuous delivery (CI/CD), rendering/transcoding, and testing.

Get more from your containers
Containers are naturally stateless and fault tolerant, making them a great fit for preemptible VMs! You save on your containerized workloads today with these affordable compute instances. Take advantage of Google Kubernetes Engine for your containerized workloads and Managed Instance Groups to painlessly and seamlessly recover from preemptions.

Enable it instantly
Simply add --preemptible to the gcloud command line and you're off to the races. There's no bidding to code for, and with per-second billing, just shut down your VMs as soon as you're done.

Question: 93

Your ed-tech start-up was originally launched in a small geography. Any user sign-ups, course progress, tests taken, etc. are captured on a self-managed MySQL database. Every user generates many such transactions. Now you're taking the application globally and preparing for a much larger influx of users from all over the world. The existing MySQL server is unlikely to be able to scale. Which convenient option can be considered?

A. Migrate to BigQuery

B. Migrate to Cloud Spanner
C. Migrate to Cloud SQL
D. Migrate to Bigtable

Answer: B

Explanation:

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

Question: 94

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

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

Answer: D

Explanation:

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

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

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

Question: 95

Which of the following is/are true about Anthos?

A. Enterprise-grade container orchestration and management service.
B. Modernizing your security for hybrid and multi-cloud deployments
C. Fully managed service mesh with built-in visibility
D. All of the Above

Answer: D

Explanation:

Anthos：

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

- Build, deploy, and optimize apps on GKE and VMs anywhere—simply, flexibly, and securely.
- Consistent development and operations experience for hybrid and multi-cloud environments.

Key features:

1. Enterprise-grade container orchestration and management service
2. Automate policy and security at scale
3. Fully managed service mesh with built-in visibility
4. Modernizing your security for hybrid and multi-cloud deployments

Question: 96

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

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

Question: 97

The government has ordered an audit of your company's dat

a. 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

Question: 98

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

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

Answer: B

Explanation:
Spanner, Cloud SQL, BigQuery
Spanner- Cloud Spanner is a fully managed, mission-critical, relational database service that offers transactional consistency at global scale, automatic, synchronous replication for high availability, and support for two SQL Google Standard SQL and PostgreSQL.
Cloud SQL- Cloud SQL is a fully-managed database service that helps you set up, maintain, manage, and administer your relational databases on Google Cloud Platform.
BigQuery- Google BigQuery is a cloud-based Architecture and provides exceptional performance as it can auto-scale up and down based on the data load and performs data analysis efficiently. On the other hand, SQL Server is based on client-server architecture and has fixed performance throughout unless the user scales it manually.

Question: 99

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

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

D. All of the Above.

Answer: D

Explanation:

Bare Metal Solution for Oracle

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

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

Question: 100

A customer has new applications to build that has to handle both batch data and streaming dat a. Which product should they choose?

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

Answer: B

Explanation:

Dataflow is the managed version of Apache Beam. Beam = Batch + Stream. Unified stream and batch data processing that's serverless, fast, and cost-effective.

Reference link- https://cloud.google.com/dataflow

Question: 101

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

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

Answer: A

Explanation:

Cloud Memorystore is an in-memory database that has sub-millisecond latency. This is ideal for caching application data that also changes once in a while.

https://cloud.google.com/memorystore

Question: 102

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.

Question: 103

Your customer has reliable information to indicate that they will use a certain amount of computing and analytics. The workloads are critical and they don't want to take a chance with VMs or BigQuery slots being unavailable during a peak period. How can they ensure that they allocate the capacity?

A. Send in the filled form to Google Cloud support to reserve the Compute Engine and BigQuery resources required.
B. Create reservations on Compute Engine and BigQuery.
C. On the day the capacity is required, set a scheduled job that will provision as many resources as required and lock it in.
D. Google Cloud is elastic for resources. You cannot reserve resources in advance; it is pay per use.

Answer: B

Explanation:

Create reservations on Compute Engine and BigQuery. You can reserve capacity in advance and use it over a period of time. You could also get a cost advantage.

=> There is no need for involved support. It is self-serve via the console.

=> You can reserve resources in advance when you have the need for it. And when you want to take a pay-per-use approach, that is also possible.

=> It is not a good idea to be lock in/hoard resources; you'll pay unnecessarily for resources. Also, it is difficult to time exactly when the demand will be.

Reference:
https://cloud.google.com/compute/docs/instances/reserving-zonal-resources
https://cloud.google.com/bigquery/docs/reservations-intro

Question: 104

An organization's applications run on an inflexible, on-premises architecture. The organization has decided to modernize their existing applications with the cloud. What may have prompted this business decision?

A. Developers want cloud providers to take full control of their application performance.
B. IT managers want cloud providers to automatically deploy their infrastructure.
C. IT managers want to stop making gradual changes.
D. Developers want to test ideas and experiment with more ease.

Answer: D

Explanation:
Modernizing applications means they can make alterations and innovate more easily.

Question: 105

An organization wants to scale their existing virtual machine architecture as quickly as possible. Why should the organization use VMware Engine?

A. To archive virtual machine instances.
B. To deploy custom APIs seamlessly.
C. To migrate virtual machines to containers.
D. To replatform virtual machines as they are.

Answer: D

Explanation:
VMware Engine helps migrate and run virtual machines in Google Cloud with minimal changes to the VM architecture.

https://cloud.google.com/learn/what-is-a-virtual-machine

Question: 106

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

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

Answer: B

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

Question: 107

What are the key features of Google Cloud Identity.

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

Answer: D

Explanation:
Cloud Identity:
A unified identity, access, app, and endpoint management (IAM/EMM) platform.
- Give users easy access to apps with single sign-on.
- Multi-factor authentication protects user and company data.
- Endpoint management enforces policies for personal and corporate devices KEY FEATURES :

Modernize IT and strengthen security

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

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

Single sign-on (SSO)

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

Works with your favorite apps

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

Question: 108

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

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

Answer: B

Explanation:

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

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

Reference:

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

Question: 109

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 the restricted.googleapis.com domain names, you'll need to create DNS records to direct traffic to the IP addresses associated with those domains.
- Your network must have appropriate routes for the destination IP ranges used by Google APIs and services. These routes must use the default internet gateway next hop. If you use the private.googleapis.com or the restricted.googleapis.com 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.

Question: 110

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

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

Answer: B

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

Question: 111

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

A. An input variable is used in making predictions. E.g. number of rooms in a house price prediction model.
B. One row of a dataset containing one or more input columns and possibly a prediction result.

C. An answer for a prediction task, either the answer produced by a machine learning system or the right answer supplied in training data. E.g. image contains a "cat".
D. The "knobs" that you tweak during successive runs of training a model. E.g. learning rate

Answer: B

Explanation:
One row of a dataset containing one or more input columns and possibly a prediction result.

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

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

Question: 112

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).

Question: 113

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

Question: 114

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

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

Answer: C

Explanation:

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

Reference:

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

Question: 115

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

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

Answer: C

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

Topic 2, Exam Pool B

Question: 116

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

A. The standby instance is used in high availability to replace the primary instance when failover occurs. The standby instance appears in the Google Cloud Console but does not get billed. When failover occurs, connections to the primary instance need to be manually transferred to the standby instance.
B. The standby instance is used in high availability to replace the primary instance when failover occurs. The standby instance appears in the Google Cloud Console but does not get billed. When failover occurs, connections to the primary instance are automatically transferred to the standby instance.
C. The standby instance is used in high availability to replace the primary instance when failover occurs. The standby instance doesn't appear in the Google Cloud Console. When failover occurs, connections to the primary instance are automati-cally transferred to the standby instance.
D. 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 ofload 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).

Question: 117

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

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

Answer: B

Explanation:
Cloud BigTable

Key features

High throughput at low latency

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

Cluster resizing without downtime

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

Flexible, automated replication to optimize any workload

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

Question: 118

In terms of Cloud SQL for MySQL Features offered by Google Cloud Platform which of the statements is/are correct?

A. Do not support Private IP (private service access).
B. Customer data is encrypted on Google's internal networks and in database tables, temporary files, and backups.
C. Do not Provide automated and on-demand backups and point-in-time recovery.
D. None of the above

Answer: B

Explanation:
Cloud SQL for MySQL:

Features

- Fully managed MySQL Community Edition databases in the cloud.
- Cloud SQL instances support MySQL 8.0, 5.7 (default), and 5.6, and provide up to 624 GB of RAM and 64 TB of data storage, with the option to automatically increase the storage size, as needed.
- Create and manage instances in the Google Cloud Console.
- Instances are available in the Americas, EU, Asia, and Australia.
- Customer data is encrypted on Google's internal networks and in database tables, temporary files, and backups.
- Support for secure external connections with the Cloud SQL Auth proxy or with the SSL/TLS protocol.
- Support for private IP (private services access).
- Data replication between multiple zones with automatic failover.
- Import and export databases using mysqldump, or import and export CSV files.
- Support for MySQL wire protocol and standard MySQL connectors.
- Automated and on-demand backups and point-in-time recovery.
- Instance cloning.

- Integration with Google Cloud's operations suite logging and monitoring.

Question: 119

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

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

Answer: B

Explanation:
Because unsecured third-party systems are a cybersecurity threat.

Question: 120

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

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

Question: 121

A large organization is struggling to manage their cloud costs effectively. They want to increase visibility 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

Question: 122

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

Question: 123

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 customer (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

Question: 124

A large travel company has thus far invested heavily in their technology team. There is strategic

pressure on the company to focus on their core business and innovate to survive in certain geographies and thrive in others. They are evaluating whether a move to Google Cloud will be good for them. Which of these reasons would be relevant for them? (choose two answer)

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

Answer: BC

Explanation:

Question: 125

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.

Question: 126

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

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

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

C. There are two types of Flavours in Apigee i.e. Apigee & Apigee Hybrid.

D. All of the above are correct.

Answer: D

Explanation:

Flavors of Apigee

Apigee comes in the following flavors:

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

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

Question: 127

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

A. Data inside cloud SQL is automatically Encrypted.

B. Cloud SQL automatically ensures your databases are reliable, secure, and scala-ble so that your business continues to run without disruption.

C. With DMS (Database Migration Service) it becomes very easy to Migration of Pro-duction Database.

D. All of the above

Answer: D

Explanation:

Cloud SQL

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

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

Key features:

Fully managed

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

Integrated

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

Reliable

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

Easy migrations to Cloud SQL

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

Question: 128

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.

Question: 129

You have contracted a partner to conduct some medical trials. This is a limited, 2-month contract. At the end of each day, you are expecting about 10 Gbs of dat
a. The data is highly sensitive. What networking option would you employ?

A. As the name indicates, set up Partner Interconnect with your partner company.
B. Setup Dedicated Interconnect with your partner.
C. Setup Cloud VPN and create an IPsec VPN tunnel with your partner.
D. Create a public IP for a VM and share that with your partners so that they can access it over the internet and share the data.

Answer: C

Explanation:
"Cloud VPN securely extends your peer network to Google's network through an IPsec VPN tunnel. Traffic is encrypted and travels between the two networks over the public internet. Cloud VPN is useful for low-volume data connections. For additional connection options, see the Hybrid Connectivity product page."

Question: 130

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."

Question: 131

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

A. On Google Cloud, create a folder corresponding to each team. Under that, there could be projects or further sub folders

as the team decides.

B. 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.

C. Since a Project could spawn other sub-Projects, on Google Cloud it is better to as-sign a folder for each Project.

D. 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

Question: 132

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 go. Recommend that the test-ers be permanently hired to plug the vulnerability issue.

C. It would be a security issue to have users come and go. Recommend that the test-ers be permanently hired to plug the vulnerability issue.

D. Create a Group with the permissions required to do the test and record their in-puts. 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.

Question: 133

What is a key difference between VMs and containers?

A. Virtual Machines take less time to launch; containers take longer to launch.

B. Virtual Machines can only run Linux; containers can run any operating system.

C. Virtual Machines use a shared operating system and are therefore lighter; containers are heavier on resources.

D. Each Virtual Machine in a machine has its own operating system; containers will share the same operating system.

Answer: D

Explanation:

VMs have their individual OSs. All containers on a node use the host operating system.

Question: 134

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 dat

a. 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 regulation. Take advantage of all the other cloud capabilities for remaining work-loads.

B. It is too risky to touch anything in such a scenario. It is best to remain entirely on-premise.

C. Regulations are guidelines. As long as the data remains encrypted, you can move it anywhere.

D. Petition the government for changes to such regulations as all industries are mov-ing to the public cloud. Then, when the regulations are eased, move to Google Cloud.

Answer: A

Explanation:

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

Question: 135

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		

Question: 136

When creating machine learning models, a key initial step is to identify the type of model required. One of these is the classification model. Which of these statements define a classification model?

A. A type of machine learning model for distinguishing among two or more discrete values. E.g. "book", "car".

B. A type of machine learning model is a meta-model maker, which classifies algo-rithms based on the quality of their output.

C. A type of machine learning model that outputs continuous (typically, floating-point) values. E.g. the predicted price of the house is $120,000.

D. A type of classic model approach that is less used today and which has been re-placed by the regression model.

Answer: A

Explanation:
A classification model classifies the incoming data into one or more discrete classes.

Question: 137

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

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

Answer: B

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

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

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

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

Regional and multi-regional configurations
No matter where your users may be, apps backed by Cloud Spanner can read and write up-to-date

strongly consistent data globally Additionally, when running a multi-region instance, your database is able to survive a regional failure, and offers industry-leading 99.999% availability.

Online schema changes with no downtime
Cloud Spanner users can make a schema change, whether it's adding a column or adding an index while serving traffic with zero downtime. Hence you now have the flexibility to adapt your database to your business needs without compromising on the

availability of your application.

Question: 138

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.

Question: 139

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

A. Use Cloud Vision API - write code to identify the text blocks, copy the data, and store it

B. Use TensorFlow - write code that will identify the type of visa and the bounding text blocks. Copy the data and then store it.

C. Use AutoML - upload other images of visas and run the model creation process which will automatically identify the visas

D. Use Data Labeling service - outsource the work of marking and extracting the in-formation to others.

Answer: A

Explanation:

Cloud Vision API allows you to programmatically identify images, text, etc. in the document. This would be the best option.

https://cloud.google.com/vision

Question: 140

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

Question: 141

You have a well established development and operations team. Your teams were managing the entire software delivery/deployment cycle on-premise. When migrating to the cloud, you want to continue 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.

Question: 142

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

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

D. All of the above.

Answer: D

Explanation:
Cloud Spanner:

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

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

Question: 143

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.

Question: 144

Which of these are defined by the following statement: a contract you have with your end customers, which, if you don't meet, you might even have to pay fines?

A. SLA - Service Level Agreement

B. SLC - Service Level Contract
C. SLO - Service Level Objective
D. SLI - Service Level Indicator

Answer: A

Explanation:

https://cloud.google.com/blog/products/devops-sre/sre-fundamentals-slis-slas-and-slos

Question: 145

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

Question: 146

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 upload any new medical images to Cloud Storage. On the basis of this statements which of the following statement is correct.

A. Create a Pub/Sub topic, and enable a Cloud Storage trigger for the Pub/Sub topic. Create an application that sends all medical images to the Pub/Sub topic.
B. Create a script that uses the gsutil command line interface to synchronize the on-premises storage with Cloud Storage. Schedule the script as a cron job.

C. In the Cloud Console, go to Cloud Storage. Upload the relevant images to the ap-propriate bucket. D. 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.

Question: 147

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 dat
a. 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.

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

Question: 148

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.

Question: 149

“With cloud messaging you can Customize and deliver messages accordingly to the predetermined time in the user's local

time zone." Comment on the above statement.

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

Answer: D

Explanation:
Firebase Cloud Messaging:

Firebase Cloud Messaging (FCM) is a cross-platform messaging solution that lets you reliably send messages at no cost.

Using FCM, you can notify a client app that new email or other data is available to sync. You can send notification messages to drive user re-engagement and retention. For use cases such as instant messaging, a message can transfer a payload of up to 4000 bytes to a client app.

Key capabilities of Firebase Cloud Messaging:

Send notification messages or data messages: Send notification messages that are displayed to your user. Or send data messages and determine completely what happens in your application code.

Versatile message targeting: Distribute messages to your client app in any of 3 ways—to single devices, to groups of devices, or to devices subscribed to topics.

Send messages from client apps: Send acknowledgments, chats, and other messages from devices back to your server over FCM's reliable and battery-efficient connection channel.

Question: 150

An e-commerce company's business has been booming. To keep up with the growth the IT team also grew. Many new people are being added and new resources are being set up. The CIO is in conversation with you over coffee one day and expresses her growing concern that they might be moving too fast. Their security checks and policies have not kept pace. She worries that somebody would make a misconfiguration or compliance violation thus exposing the company to data and privacy loss. What can you advise her?

A. Use Cloud Identity-Aware Proxy to allow only specific users to access the data.
B. Use Security Command Center to have a centralized view of assets and get noti-fied on misconfigurations and vulnerabilities.
C. Use Cloud Data Loss Prevention to prevent the loss of any data.
D. Use Cloud Armor to block any DDoS attacks that could be a threat.

Answer: B

Explanation:
Security Command Center is the right tool for this use case. It can check resources for security issues and notify you when issues are found.

https://cloud.google.com/security-command-center

Question: 151

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

A. Use the popular open-source libraries SciPy and NumPy to create machine learn-ing models.
B. Use the Unified AI Platform to create a custom TensorFlow model.
C. Use BigQuery ML to create machine learning models using SQL queries.

D. Integrate the Cloud Vision API and the Cloud Speech API to create a custom mod-el that will suit the retail sector.

Answer: C

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

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

Question: 152

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 preallocated. 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 VMs. Get an estimate of usage. Then plan for a VM capacity of 25% above the load test value.
B. Use the Managed Instance Group with Compute Engine

C. Inform them that new-age companies are using microservices, containers, and Kubernetes for this and they can plan to rewrite the app quickly.
D. Inform them that using a serverless option will take care of the scaling and they can move to Cloud Run or App Engine.

Answer: B

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

Instance groups □ Send feedback

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

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

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

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

Question: 153

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 â€Billing Account Userâ€ 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

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

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

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

Question: 154

You are working with a user to set up an application in a new VPC behind a firewall and it is no-ticed that the user is concerned about data egress. Therefore, to provide assistance you want to con-figure the fewest open egress ports. Which of the following statement is correct?

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

Answer: B

Explanation:
Implied rules Every VPC network has two implied firewall rules. These rules exist, but are not shown in the Cloud Console: Implied allow egress rule. An egress rule whose action is allow, destination is 0.0.0.0/0, and priority is the lowest possible (65535) lets any instance send traffic to any destination, except for traffic blocked by Google Cloud. A higher priority firewall rule may restrict outbound access. Internet access is allowed if no other firewall rules deny outbound traffic and if the instance has an external IP address or uses a Cloud NAT instance. For more information, see Internet access requirements.

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

Question: 155

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

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

Answer: E

Explanation:

Question: 156

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

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

Answer: A

Explanation:

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

Question: 157

Google offers Firebase, In terms of Firebase Console, any particular message that has to be delivered to a customer at a certain degree of change in behavior can be managed through

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

Answer: B

Explanation:
You can send notification messages using the Notifications composer in the Firebase console. Though this does not provide the same flexibility or scalability as sending messages with the Admin SDK or

the HTTP and XMPP protocols, it can be very useful for testing or for highly targeted marketing and user engagement. The Firebase console provides analytics-based A/B testing to help refine and improve marketing messages.
After you have developed logic in your app to receive messages, you can allow non-technical users to send messages per the instructions on the Notifications page in the Firebase Help Center.

Question: 158

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

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

Answer: C

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

Question: 159

One of your customers used to have a private data center. While within their data center itself, they were consuming some Google services via API calls and other public, well-known addresses published by Google. Now they're evacuating their private data center and are moving to Google Cloud. Could they improve some of their existing architecture with respect to security?

A. Use VPC Peering with the Google Cloud organization so that you can directly use services using only private IPs.
B. Use private addresses only. No additional configuration is required. All Google services will be accessible within Google Cloud on private addresses.
C. Use Shared VPCs with the Google Cloud organization so that you can directly use services using only private IPs.
D. Enable Private Google Access so that they can remove public IP addresses.

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. f 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."

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

Question: 160

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

A. Teamwork over individual work

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

Answer: D

Explanation:

Question: 161

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

Question: 162

Firebase Hosting provides the following services-

A. Dynamic content
B. Static content.
C. Microservices.
D. All of the Above.

Answer: D

Explanation:
Firebase Hosting- Firebase Hosting provides fast and secure hosting for your web app, static and dynamic content, and microservices.

Firebase Hosting is production-grade web content hosting for developers. With a single command, you can quickly deploy web apps and serve both static and dynamic content to a global CDN (content delivery network). You can also pair Firebase Hosting with Cloud Functions or Cloud Run to build and host microservices on Firebase.

Key capabilities of Firebase Hosting:

Serve content over a secure connection:- The modern web is secure. Zero-configuration SSL is built into Firebase Hosting, so content is always delivered securely.

Host static and dynamic content plus microservices:- Firebase Hosting supports all kinds of content for hosting, from your CSS and HTML files to your Express.js microservices or APIs.

Deliver content fast: Each file that you upload is cached on SSDs at CDN edges around the world and served as gzip or Brotli. We auto-select the best compression method for your content. No matter where your users are, the content is delivered fast.

Question: 163

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 storage. Retain the data in use in a dual-region bucket.
B. Retain the data in use in a single region bucket with standard storage.
C. Retain the data in use in a multi-region bucket.
D. 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'	Typical monthly availability
multi-region	99.95%	>99 99%

dual-region	99.95%	>99.99%

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

Question: 164

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

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

Answer: C

Explanation:

Firebase Crashlytics:

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

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

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

Question: 165

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

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

Answer: D

Explanation:

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

Question: 166

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 donttake full advantage of cloud platform features, such as horizontal scalability, finegrained pricing, and highly managed services

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

Question: 167

Your client's IT environment has so far been on-premises. They run a mix of applications and databases on Linux and Windows. They want to move to Google Cloud in the easiest manner possi-ble. What are their best options?

A. Compute Engine with VMs with either Linux or Windows OS.

B. App Engine Standard

C. Cloud Functions

D. Cloud Run

Answer: A

Explanation:

Compute Engine allows you to allocate VMs with different OSs - Windows and Linux, included.

Question: 168

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.

Question: 169

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 possible. If yes, there will be a significantly higher charge be-cause the effort is significantly higher too.

B. Yes, that is possible, but there will be an additional charge of 9% for the service because that is the additional effort required.

C. Yes, that is possible. There is hardly any difference to provide another 0.009% availability.

D. Ask them for the reasonable downtime they are willing to absorb. 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.

E. 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

Question: 170

An organization wants to evaluate the performance of their entire cloud infrastructure, including metrics like server uptime and response rate reports. Which Google Cloud tool should the organization use?

A. Cloud Trace
B. Cloud Monitoring
C. Cloud Profiler
D. Cloud Debugger

Answer: B

Explanation:
Because Cloud Monitoring enables users to monitor the performance of their entire cloud infrastructure.

Question: 171

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

Question: 172

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

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

Question: 173

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

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

Answer: B

Explanation:
Coldline storage is a very low cost highly durable storage service for data archiving, online backup, and disaster recovery. Coldline storage is the best choice for data that you plan to access at most once a year due to its slightly lower availability, 90 day minimum storage duration cost for data access, and higher per operation costs.
Nearline and Coldline are for backup and archival storage and having the highest availability for both with 99.9 percent.

Question: 174

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

Question: 175

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

Explanation:

Question: 176

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 happening quickly thereafter from these new accounts. Different parts of the application are being explored, 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 system. Use Cloud Asset Inventory to see if there have been changes in the inventory.
B. Bots are creating accounts and then using them. Use Google Cloud's Web App and API Protection (WAAP).
C. Bots are creating accounts and then using them. Use Identity-Aware Proxy to re-strict the users to known users.
D. Automated testing tools might still be running and creating accounts. 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

Question: 177

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

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

Answer: A

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

Question: 178

Your customer's IT team is in the process of modernizing their customer-facing applications. They've witnessed others getting good results from employing microservices, and they're keen to adopt it themselves. The first application that they are modernizing has about 5 different sub-parts, which they have identified will be the services. They also identify that each of them has different scale requirements - some services like user login are less frequently used while others like transac-tions are heavily used. What technical strategy would you recommend for them?

A. Containerize the services and orchestrate them with Google Kubernetes Engine.
B. Retain the original application in Compute Engine and scale it as needed using Managed Instance Groups.
C. Retain the original application as a backup and also for separately scaling the ser-vices, create new application binaries.
D. Retain the original application in Compute Engine and scale it as needed using Unmanaged Instance Groups.

Answer: A

Explanation:
Containers and Kubernetes are ideal for the kind of requirement mentioned here - separate microservices that need to scale independently.
Google Kubernetes Engine (GKE) provides a managed environment for deploying, managing, and

scaling your containerized applications using Google infrastructure. The GKE environment consists of multiple machines (specifically, Compute Engine instances) grouped together to form a cluster.

Reference link- https://cloud.google.com/kubernetes-engine/docs/concepts/kubernetes-engine- overview

Question: 179

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.

Question: 180

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

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

Answer: D

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

With Bare Metal Solution, Google Cloud provides and manages the core infrastructure, the network,

the physical and network security, and hardware monitoring capabilities in an environment from which you can access all of the Google Cloud services. The core infrastructure includes secure, controlled-environment facilities, and power.
The Bare Metal Solution also includes the provisioning and maintenance of custom, sole-tenancy servers with local SAN, and smart hands support.
The network, which is managed by Google Cloud, includes a low-latency Partner Interconnect connection into the customer Bare Metal Solution environment.
The available Google Cloud services include private API access, management tools, support, and billing.

Question: 181

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

A. Bare Metal Solutions

B. Preemptible Virtual Machines.
C. Google Cloud VM Instances
D. None of the above.

Answer: B

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

Question: 182
In terms of Infrastructure as a Service (IaaS) what are the benefits of it?

A. IaaS offers virtually infinite flexibility and scalability, enterprises can get their work done more efficiently, ensuring faster development life cycles.
B. IaaS resources are regularly available to businesses when they need them. As a result, enterprises reduce delays when expanding infrastructure and, alternatively, don't waste resources by overbuilding capacity.
C. IaaS resources are used on demand and enterprises only have to pay for the compute, storage, and networking resources that are actually used, IaaS costs are fairly predictable and can be easily contained and budgeted for.
D. All of the Above

Answer: D

Explanation:
These are the feature of Infrastructure as a Service (IaaS)
It's economical
Because IaaS resources are used on demand and enterprises only have to pay for the compute, storage, and networking resources that are actually used, IaaS costs are fairly predictable and can be easily contained and budgeted for.
It's efficient
IaaS resources are regularly available to businesses when they need them. As a result, enterprises reduce delays when expanding infrastructure and, alternatively, don't waste resources by overbuilding capacity.
It boosts productivity
Because the cloud provider is responsible for setting up and maintaining the underlying physical infrastructure, enterprise IT

departments save time and money and can redirect resources to more strategic activities.

It's reliable

IaaS has no single point of failure. Even if any one component of the hardware resources fails, the service will usually still remain available.

It's scalable

One of the biggest advantages of IaaS in cloud computing is the capability to scale the resources up and down rapidly according to the needs of the enterprise.

It drives faster time to market

Because IaaS offers virtually infinite flexibility and scalability, enterprises can get their work done more efficiently, ensuring faster development life cycles.

Question: 183

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

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

Answer: D

Explanation:

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

Question: 184

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

Question: 185

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

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

Answer: D

Explanation:
Preemptible instances function like normal instances but have the following limitations:
- > Compute Engine might stop preemptible instances at any time due to system events. The probability that Compute Engine will stop a preemptible instance for a system event is generally low, but might vary from day to day and from zone to zone depending on current conditions.

- > Compute Engine always stops preemptible instances after they run for 24 hours. Certain actions reset this 24-hour counter.

- > Preemptible instances are finite Compute Engine resources, so they might not always be available.

- > Preemptible instances can't live migrate to a regular VM instance, or be set to automatically restart when there is a maintenance event.

- > Due to the above limitations, preemptible instances are not covered by any Service Level Agreement (and, for clarity, are excluded from the Compute Engine SLA).

- > The Google Cloud Free Tier credits for Compute Engine do not apply to preemptible instances.

Question: 186

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

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

Answer: A

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

Question: 187

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.

Question: 188

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

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

Answer: C

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

Question: 189

Your customer has a reporting tool that is only occasionally used by the leadership team. Usage of it is frequent - once a week, once a month, or once the quarter. They want to run this application in a cost-effective manner. What are the compute options available on Google Cloud which would be suitable? (Choose Two answer)

A. Cloud Run
B. Cloud App Engine Standard
C. Compute Engine
D. Kubernetes Engine

Answer: AB

Explanation:
Since the use of the tool is infrequent/intermittent, you can choose to compute options that are serverless. Both Cloud Run and Cloud App Engine Standard are serverless options that can shut down to zero. Since cost-effectiveness is a requirement, this will not cost anything during the periods it is

not used.

Question: 190

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

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

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

Answer: C

Explanation:
The above is a real world example of a company named Planet, where they sent around 80+ satellites to take pictures of earth every day, 24 hours. They run around 40,000 preemptible VMs concurrently.

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

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

Question: 191

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.

Question: 192

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

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

Answer: D

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

After you create a snapshot of a boot persistent disk, you can apply data from that snapshot to new persistent disks. It is only possible to apply data from a snapshot when you first create a persistent disk. You cannot apply a snapshot to an existing

persistent disk, or apply a snapshot to persistent disks that belong to a different project than that snapshot.

Question: 193

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.

Question: 194

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

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

Question: 195

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

A. Kubernetes uses Docker to deploy, manage, and scale containerized applications.
B. Difference between Docker and Kubernetes relates to the role each play in con-tainerizing and

running your applications

C. Kubernetes can be used with or without Docker.

D. All of the above.

Answer: D

Explanation:

Kubernetes vs. Docker

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

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

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

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

Question: 196

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.

Question: 197

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

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

Answer: A

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

- > Sharing a single project will cause a conflict of resources, billing, concerns, etc.
- > One folder per project is unnecessary overuse of abstraction/grouping.
- > Teams and projects in a company should ideally be centrally managed in a single Organization.

Question: 198

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

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

Answer: D

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

Question: 199

The government has mandated that companies in a particular section of healthcare must retain all
the data they collect for a period of 10 years in case an audit needs to be done. Your client, who is in that industry, needs to follow regulations. In addition, your client wants to do an analysis of the data quite frequently in the first year. They also don't want to be liable for any data beyond year 10. What would recommend for your customer?

A. Use Cloud Storage with nearline storage in year one and Coldline storage thereaf-ter. Use Object lifecycle management to move between storage types and delete them after 10 years.
B. Use Cloud Storage with standard storage in year one and Coldline storage there-after. Set a Cloud Scheduler trigger for 1 year to change storage types and 10 years to delete the data.
C. Use Cloud Storage with standard storage in year one and archival storage thereaf-ter. Use Object lifecycle management to move

between storage types and delete them after 10 years.

D. Use Cloud Storage with standard storage in year one and Coldline storage there-after. Set a Cloud Tasks to trigger for 1 year to change storage types and 10 years to delete the data.

Answer: C

Explanation:
Cloud storage supports Object Lifecycle Management. To support common use cases like setting a Time to Live (TTL) for objects, retaining noncurrent versions of objects, or "downgrading" storage classes of objects to help manage costs, Cloud Storage offers the Object Lifecycle Management feature.

Standard storage is recommended for frequently accessed data and Archive for data accessed less than once a year.

Nearline, Coldline, and Archive offer ultra-low-cost, highly-durable, highly available archival storage. For data accessed less than once a year, Archive is a cost-effective storage option for the long-term preservation of data. Coldline is also ideal for cold storage—data your business expects to touch less than once a quarter. For warmer storage, choose Nearline: data you expect to access less than once a month, but possibly multiple times throughout the year.

Question: 200

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

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

Answer: B

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

Question: 201

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.

Question: 202

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

Question: 203

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

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

Answer: D

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

load balancer can be in Premium Tier or Standard Tier.

-> Google Cloud Armor security policies and IP DENY lists and ALLOW lists are available only for HTTP(S) load balancing.
Reference link- https://cloud.google.com/armor/docs/security-policy-overview

Question: 204

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

Question: 205

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

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

Answer: B

Explanation:
Enable Services and APIs

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

Question: 206

Cloud Data Loss Prevention (DLP) is a fully managed service designed to help discover, classify, and protect the most sensitive dat

a. DLP provides three key features (Select Three Answers)

A. Classification
B. De-identification
C. De-classification
D. Inspection
E. Reinspection

Answer: ABD

Explanation:
Classification. De-classification and Inspection
Classification is the process to inspect the data and know what data we have, how sensitive it is, and the likelihood. Inspection and classification happen here.
De-identification is the process of removing, masking, replacing information from data.

Reference link- https://cloud.google.com/dlp/docs

Question: 207

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 name. Look for similar bucket names that have a "." in it.
B. Previously deleted bucket names in the same project cannot be reused. There must have been an older bucket with the same name.
C. Bucket names in Cloud storage have to be globally unique
D. 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.

Question: 208

Your company has signed up with a cloud provider and you will be using storage and vir-tual machines with the provider. The provider has provided your organization some expec-tations for what the service should perform at. What type of agreement provides a guar-antee of a certain level of service such as “Uptime”?

A. Performance Agreement
B. Interconnection Agreement
C. Warranty
D. Service Level Agreement

Answer: D

Explanation:
Service Level Agreement (SLA)

A service level agreement (SLA) is a contract between a service provider (either internal or external) and the end user that defines the level of service expected from the service provider. Some common SLA's are uptime, Response Time, etc.

Question: 209

Which of the following are the current options for paid support in GCP? (Select Three An-swer)

A. Premier
B. Standard
C. Enhanced
D. Role
E. Premium

Answer: BCE

Explanation:
Because GCP provides three options for paid support which are Standard, Enhanced and Premium. Basic Support is included with your Google Cloud subscription which cover only Case, phone, and chat support for billing issues only

Reference link- https://cloud.google.com/support

Question: 210

What cloud service model would you want to select if you want to solve a particular busi-ness problem by providing CRM services in the cloud to your enterprises?

A. CaaS
B. SaaS
C. PaaS
D. IaaS

Answer: B

Explanation:
SaaS – Software as a Service (SaaS) provides you a complete product that is run and managed by the service provider. You worry

only about using the software and not about infrastructure.
SaaS provides the lowest level of flexibility and management control over the infrastructure.
(Example: Google Gsuite and MS O365)

Question: 211

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

A. Supports autoscaling

B. Supports Docker containers

C. Supports mainframe connectivity

D. Source code is written in specific versions of the supported programming lan-guages only

Answer: AB

Explanation:

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

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

Question: 212

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

A. About 12 days.
B. About 2 years.
C. About 100 days.

D. About 24 hours.

Answer: C

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

number by 8.
Total data transferred (in bits) = 100 * 1,000,000,000,000 * 8 bits
Speed = 100Mbps = 100 * 1,000,000. i.e. 100 million bits are transferred per second.
Hence time taken to transfer all the data = Total Data/Speed = 8,000,000 seconds.
Number of seconds in a day = 24*60*60 = 86,400
Total time taken in days = 8,000,000/86,400 = 92.59 days

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

Question: 213

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

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

Answer: B

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

Question: 214

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

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

Answer: B

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

GKE resources.

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

Question: 215

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.

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

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

Question: 216

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

Question: 217

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

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

Answer: B

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

- > Pub/Sub allows services to communicate asynchronously, with latencies on the order of 100 milliseconds.
- > Pub/Sub is used for streaming analytics and data integration pipelines to ingest and distribute data. It is equally effective as messaging-oriented middleware for service integration or as a queue to parallelize tasks.
- > Pub/Sub enables you to create systems of event producers and consumers, called publishers and subscribers. Publishers communicate with subscribers asynchronously by broadcasting events, rather than by synchronous remote procedure calls (RPCs).

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

Question: 218

A customer is migrating there on-promises data analytics solution to Google Cloud. The current solution has a lot of data being read form and written to disk. The performance of this approach has occasionally been a bottleneck for a scale of operations that your cus-tomer has. The application is fault tolerant and can with stand machine going down fre-quently. In moving to Google Cloud they are asking your advice on any way to improve performance?

A. Use Big Query Which has very fast data access and analysis
B. Use Cloud Storage which can be central, scalable storage
C. Use local SSDs with the VMs
D. Use Persistent Disk with the VMs

Answer: C

Explanation:
Local SSDs are attached to the VM and have very high throughput. However, when the VM shuts down, The local SSD is also shut down, Since our Workload here is foult tolerant, than is not an issue.

Question: 219

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

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

Answer: CDE

Explanation:
Sustained, Committed and Preemptible

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

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

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

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

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

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

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

Reference link– https://cloud.google.com/compute/docs/instances/preemptible

Question: 220

Your customer is moving from AWS to Google Cloud. Data also needs to be moved. There is about 50TB of dat a. On AWS, the data resides in an S3 bucket. It is going to be moved to Cloud Storage. Data is also being continuously generated on S3 prior to the cutover. It is preferable that this is also periodically transferred. What is the best way to move the data?

A. Use the gsutil command-line option
B. Use the Google Cloud console to drag and drop the files easily
C. Use the Storage Transfer Service
D. Use a Transfer Appliance

Answer: C

Explanation:
Storage Transfer Service provides options that make data transfers and synchronization easier. We can also schedule one-time transfer operations or recurring transfer operations.

Reference link- https://cloud.google.com/storage-transfer/docs/overview

Reference link- https://cloud.google.com/architecture/transferring-data-from-amazon-s3-to-cloud-storage-using-vpc-service-controls-and-storage-transfer-service

Question: 221

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

Question: 222

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

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

Answer: B

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

Question: 223

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

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

Answer: D

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

Topic 3, Exam pool (Latest)

Question: 224

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

Explanation:

Question: 225

An organization wants to build an entirely new infrastructure and applications in the cloud.
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: C

Explanation:
A Greenfield approach is a brand-new implementation , where companies then add their needed configurations and customizations. This approach provides a clean slate to start from, does not carry over needless customizations and technical debt, and provides a solid foundation for business process re-engineering.
A greenfield deployment is the design, installation and configuration of computer infrastructure where none existed before, for example, in a new office. In contrast, a brownfield deployment is an upgrade or addition to existing infrastructure using legacy components.

Question: 226

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

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

Answer: B

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

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

Question: 227

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.

Question: 228

An organization finds that the amount of cash in their vending machines doesn't match the value of items sold. They have decided to upgrade their vending machines with cloud-based mobile payment systems.
How could the organization benefit from this upgrade?

A. They could relax data access permissions.
B. They could reduce their error budget overspend.
C. They could improve their perimeter security.
D. They could view data history to see transactions.

Answer: D

Explanation:

Question: 229

An organization's developers are growing increasingly frustrated by the limitations of their onpremises 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.)

Question: 230

An online retail organization wants to optimize their service.
What is an example of unstructured data that they can use to make decisions?

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

Answer: A

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

Question: 231

An organization wants to add a new function to their application. They want to write the code and let the public cloud provider handle the infrastructure.
Which infrastructure solution should they use?

A. Virtual machines
B. Bare Metal Solution
C. Serverless computing
D. Container Registry

Answer: C

Explanation:
Serverless computing , as public cloud prouder(eg. google) will mange the infra things

Question: 232

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

Explanation:

Question: 233

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.

Question: 234

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

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

Answer: C

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

Question: 235

An organization wants full control of their virtual machine infrastructure for a custom home-grown application with a product that autoscales and automatically updates.
Which Google Cloud product or solution should the organization use?

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

Answer: C

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

Question: 236

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

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

Answer: D

Explanation:

Question: 237

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

Explanation:

Question: 238

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

Explanation:

Question: 239

What is an example of unstructured data that organizations can capture from social media?

A. Post comments
B. Tagging
C. Profile picture
D. Location

Answer: A

Explanation:
https://treehousetechgroup.com/8-examples-of-unstructured-data/

Question: 240

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

Question: 241

An organization provides a loyalty program for its customers. It recently partnered with other businesses so that customers can get loyalty points at a range of other stores.
Why should the organization use application programming interfaces (APIs)?

A. To migrate all partner data for disaster recovery
B. To analyze and publish loyalty program statistics to a dashboard
C. To personalize recommendations for loyalty card users
D. To connect third-party systems to ensure up-to-date information

Answer: D

Explanation:

Question: 242

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

Explanation:

Question: 243

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

A. Cloud Monitoring
B. Cloud Trace
C. Cloud Logging
D. Cloud Debugger

Answer: A

Explanation:
https://cloud.google.com/monitoring

Question: 244

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

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

Answer: A

Explanation:
https://cloud.google.com/vision

Question: 245

An organization wants to use multiple marketing datasets to forecast user acquisition.
How should they use cloud technology to gain new insights from the data?

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

Answer: D

Explanation:

Question: 246

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

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

Answer: C

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

Question: 247

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

Explanation:

Question: 248

An organization has decided to modernize their applications in the cloud to keep up with their customers' needs. What may have prompted this business decision?

A. Their on-premises applications only autoscale to meet demand.
B. They want to change from a pay-as-you-go model to a capital expenditure model.
C. Their source code changes erroneously without developer interaction.
D. Their on-premises applications take months to update and deploy.

Answer: D

Explanation:

Question: 249

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

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

Answer: D

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

Question: 250

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

Explanation:

Question: 251

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

Explanation:

Question: 252

An organization has servers running mission-critical workloads on-premises around the world. They want to modernize their infrastructure with a multi-cloud architecture.
What benefit could the organization experience?

A. Ability to disable regional network connectivity during cyber attacks
B. Ability to keep backups of their data on-premises in case of failure
C. Full management access to their regional infrastructure
D. Reduced likelihood of system failure during high demand events

Answer: D

Explanation:

Question: 253

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

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

Answer: B

Explanation:

Question: 254

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.

Question: 255

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

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

Answer: D

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

Question: 256

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.

Question: 257

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

Explanation:

Question: 258

What does Cloud Debugger help an organization do?

A. Implement code updates in real time without affecting the service level objective (SLO).
B. Inspect source code in real time without affecting user downtime.
C. Manage code and accelerate application development.
D. Analyze live source code during user downtime.

Answer: B

Explanation:
Cloud Debugger is a feature of Google Cloud Platform that lets you inspect the state of an application, at any code location, without stopping or slowing down the running app. Cloud Debugger makes it easier to view the application state without adding logging statements.

Question: 259

Why is data stored in Google Cloud secure and private?

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

Answer: C

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

Question: 260

An organization needs to categorize text-based customer reviews on their website using a pretrained machine learning

model.
Which Google Cloud product or service should the organization use?

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

Answer: A

Explanation:

https://cloud.google.com/natural-language
Use entity analysis to find and label fields within a document—including emails, chat, and social media—and then sentiment analysis to understand customer opinions to find actionable product and UX insights.

Question: 261

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

Explanation:

Question: 262

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

Question: 263

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

Explanation:

Question: 264

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

Explanation:

Question: 265

An organization is altering their gaming product so that it is compatible with cloud technology.
What can they expect when moving from traditional technology to cloud technology?

A. No change to existing responsibilities
B. A shift toward OpEx
C. A shift toward using structured data
D. Increased hardware maintenance

Answer: B

Explanation:

Question: 266

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

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

Answer: C

Explanation:

Question: 267

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

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

Answer: B

Explanation:

Question: 268

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

A. Customer expectations can be adjusted without using marketing tools
B. Seamless changes can be made without causing any application downtime.

C. GKE handles version control seamlessly and out of the box D. GKE is well suited for all monolithic applications

Answer: B

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

Question: 269

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

Explanation:

Question: 270

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

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

Answer: A

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

Question: 271

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.

Question: 272

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

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

Answer: B

Explanation:

https://www.crowdstrike.com/cybersecurity-101/zero-trust-security/
Zero Trust is a security framework requiring all users, whether in or outside the organization's network, to be authenticated, authorized, and continuously validated for security configuration and posture before being granted or keeping access to applications and data

Question: 273

An organization wants to search hundreds of scanned documents for key information like dates, names, and other specific words.
Why should the organization use application programming interfaces (APIs)?

A. To replace the scanned documents with an online survey
B. To ingest data in real time and encrypt unmatched words
C. To create digital versions of the documents and locate key information D. To transform the documents

into unstructured data.

Answer: D

Explanation:
The text from the PDF/scanned documents/images gets converted into JSON (unstructured file) which will be further used for search.

Question: 274

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

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

Answer: B

Explanation:

Question: 275

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

Explanation:

Question: 276

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

Explanation:

Question: 277

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

Explanation:

Question: 278

Which policy helps Google Cloud keep customer data private?

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

Answer: B

Explanation:

Question: 279

Why should an organization consider the total cost of ownership (TCO) when moving from onpremises 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

Explanation:

Question: 280

What is monitoring within the context of cloud operations?

A. Observing cloud expenditure in real time to ensure that budgets are not exceeded
B. Collecting predefined and custom metrics from applications and infrastructure
C. Tracking user activities to guarantee compliance with privacy regulations
D. Tracing user location to document regional access and utilization

Answer: B

Explanation:

Question: 281

An organization notices that some of their cloud expenditures are too high.
What should the organization do to control costs?

A. Streamline the hardware procurement process to reduce costs.
B. Share cost views with the departments to establish more accountability.
C. Change the cost model from operational expenditure to capital expenditure.
D. Ensure that all could resources are tagged with a single tag.

Answer: B

Explanation:

Question: 282

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

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

Answer: B

Explanation:

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

Question: 283

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.

Question: 284

How would an organization benefit from using Looker?

A. Optimal identity and access management
B. Leading serverless warehousing technology
C. Robust data roll-back accuracy
D. Advanced business intelligence and analytics

Answer: D

Explanation:

Looker is a business intelligence software and big data analytics platform that helps you explore, analyze and share real-time business analytics easily.

Question: 285

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.

Question: 286

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

Explanation:

Question: 287

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

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

Answer: B

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

Question: 288

How does Cloud SQL help organizations create business insights?

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

Answer: A

Explanation:

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

Question: 289

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

Question: 290

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

Question: 291

How can a streaming service meet global compliance requirements using the cloud?

A. By automatically encrypting personally identifiable information
B. By obtaining a business license to operate in a new market
C. By allowing users to disable two-factor authentication
D. By adhering only to data polices of the country in which the head office is registered

Answer: A

Explanation:

Question: 292

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

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

Answer: C

Explanation:

Question: 293

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

Question: 294

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

A. To create personalized product recommendations for customers

B. To optimize their on-premises legacy system stability
C. By manually linking each inventory system to the website on a case-by-case basis
D. By programmatically connecting the inventory system to their website

Answer: D

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

Question: 295

An organization is making a strategic change to customer support in response to feedback. They plan to extend their helpline availability hours.
Why is the organization making this change?

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

Answer: C

Explanation:

Question: 296

Which technology allows organizations to run multiple computer operating systems on a single piece of physical hardware?

A. Hypervisor
B. Containers
C. Serverless computing
D. Open source

Answer: A

Explanation:

Question: 297

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

application?

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

Answer: B

Explanation:

Question: 298

An organization wants to digitize and share large volumes of historical text and images. Why is a public cloud a better option than an on-premises solution?

A. In-house hardware management
B. Provides physical encryption key
C. Cost-effective at scale
D. Optimizes capital expenditure

Answer: C

Explanation:

Question: 299

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

Explanation:

Question: 300

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

Explanation:

Question: 301

An organization's web developers and operations personnel use different systems.
How will increasing communication between the teams reduce issues caused by silos?

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

Answer: D

Explanation:

Question: 302

An e-commerce organization is reviewing their cloud data storage.
What type of raw data can they store in a relational database without any processing?

A. Product inventory
B. Product photographs
C. Instructional videos
D. Customer chat history

Answer: A

Explanation:

Question: 303

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

Explanation:

Question: 304

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

Explanation:

Question: 305

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

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

Answer: B

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

Question: 306

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

Question: 307

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

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

Answer: A

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

Question: 308

An organization decides to migrate their on-prenwses environment to the cloud They need to determine which resource components still need to be assigned ownership
Which two functions does a public cloud provider own? (Choose 2)

Choose 2 answers

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

Answer: CD

Explanation:

Question: 309

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.

Question: 310

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

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

Answer: C

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

Question: 311

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

Explanation:

Question: 312

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

Explanation:

Question: 313

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

Explanation:

Question: 314

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

Explanation:

Question: 315

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

A. Maintaining overall system operability
B. Maintaining customer-facing content

C. Monitoring data center servers
D. Monitoring computer networks

Answer: B

Explanation:

Question: 316

An organization is looking for a business intelligence solution that allows individual employees and end users to analyze business data and generate insights.
Which Google Cloud product or service should the organization use?

A. Looker
B. Cloud Spanner
C. BigQuery
D. Dataflow

Answer: A

Explanation:

Question: 317

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

Explanation:

Question: 318

An organization needs frequent access to only a subset of their dat
a. They want to reduce costs by depositing the rest of their data across Nearline Coldlme and Archive repositories
Which Google Cloud product should the organization use?
A. Filestore
B. Cloud Spanner

C. Data Catalog
D. Cloud Storage

Answer: D

Explanation:
Per Google docs, specifically for GCP Cloud Storage there exists four types of storage with one of them, standard storage, being described as "storage for data that is frequently accessed ("hot" data) and/or stored for only brief periods of time." https://cloud.google.com/storage

Question: 319

A retail organization has moved all of their inventory data to a relational database in the cloud. What functionality does a relational database offer?

A. It analyzes unstructured data which can then be accessed in multiple regions
B. It stores transactional data which can then be accessed electronically
C. It stores large amounts of raw data in its original format
D. It rapidly analyzes large and multi-dimensional datasets

Answer: B

Explanation:
A relational database offers the functionality of storing transactional data, which can then be accessed electronically. Relational databases store structured data that can be organized in tables with defined relationships between them. This makes them well-suited for transactional data, such as inventory data, that needs to be accessed and updated frequently.

Question: 320

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

Explanation:
To follow Site Reliability Engineering (SRE) principles, the recommended approach when introducing a new system, like an image recognition login system, is to minimize risk and test in a controlled environment. SRE emphasizes progressive rollouts and monitoring to ensure reliability, stability, and security.
Option A: Roll out the new system to a subset of employees to test it out is the correct answer because this aligns with SRE practices, such as:
Canary Releases: Deploying a new feature to a small group of users (subset of employees) allows the organization to test the system in a real-world environment with minimal risk. This practice helps identify any

potential issues, gather feedback, and monitor system behavior without impacting the entire organization.
Gradual Rollouts: SRE encourages gradual rollouts to detect and mitigate any failures early, reducing the blast radius in case of a malfunction or security issue. This ensures the reliability of the system while maintaining service quality.
Monitoring and Observability: Rolling out to a subset allows for comprehensive monitoring and collection of metrics to ensure the system performs as expected. If issues arise, they can be quickly identified and resolved before a wider deployment.
Option B: Roll out the new system to all employees to collect as much data as possible is not advisable under SRE principles because it poses a higher risk of widespread failure or security issues. Option C: Avoid rolling out the new system because it may have security flaws and Option D: Avoid rolling out the new system because it may violate privacy policy are also incorrect in the context of SRE. While security and privacy are crucial, outright avoidance does not align with SRE practices. Instead, risk is managed through controlled rollouts, testing, and monitoring.
Reference:
Google Cloud SRE Workbook: Canarying Releases
Google Cloud SRE Principles: Monitoring, Progressive Rollouts
Google Cloud Architect's Guide: Reliability and Risk Management

Question: 321

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

Explanation:
The correct answer is B. Machine learning. Here's why:
Context of the Questio n : The question describes an application that can diagnose medical conditions based on images submitted by users. This is a task typically handled by machine learning (ML) models, specifically those focused on image recognition and classification.
Google Cloud Product Relevance:
Machine learning (ML) is the correct service because it refers to the use of Google Cloud's AI/ML capabilities, which include services such as AutoML Vision and Vertex AI. These tools can build and deploy machine learning models specifically for image recognition tasks, such as identifying medical conditions from images.
AutoML Vision is designed to create custom ML models with minimal ML expertise, making it suitable for diagnosing medical conditions from images.
Vertex AI allows for training, evaluating, and deploying ML models, including complex models for medical image diagnostics.
Why Not Other Options:
A . App Engine: App Engine is a platform-as-a-service (PaaS) used for building and deploying web applications, not specifically for ML or image recognition tasks.

C . Cloud Logging: Cloud Logging is used to store, search, analyze, and monitor log data and events. It does not have any capabilities related to diagnosing medical conditions or image recognition.
D . Cloud Profiler: Cloud Profiler is a performance analysis tool for applications running in production. It helps identify bottlenecks and optimize resource usage but has no relevance to image diagnosis or machine learning.
Google Cloud Digital Leader Reference:
For more information on Machine Learning with Google Cloud, refer to the Google Cloud documentation on AutoML Vision and Vertex AI.
Understanding the basic principles and applications of Google Cloud's ML services is crucial for the Google Cloud Digital Leader exam, which emphasizes knowing how different services solve business challenges.

Question: 322

An information security organization must ensure that its service providers have ISO 27001 certification. They must also access supporting documentation. Which Google Cloud tool should they use?

A. Network Intelligence Center
B. Cloud Monitoring
C. Compliance Reports Manager
D. Security Command Center

Answer: C

Explanation:
The correct answer is C. Compliance Reports Manager. Here's why:
Context of the Questio n : The organization needs to ensure that its service providers have ISO 27001 certification and access supporting documentation. ISO 27001 is an international standard for information security management systems (ISMS), and organizations often need to verify compliance through certified reports.
Google Cloud Product Relevance:
Compliance Reports Manager is a tool within the Google Cloud Console that provides access to various compliance and security documentation, including ISO 27001 certification. It allows customers to view and download compliance reports, making it the appropriate choice for verifying certifications and accessing related documents.
Why Not Other Options:
A . Network Intelligence Center: This tool is designed for network monitoring, visibility, and troubleshooting and is unrelated to compliance management.
B Cloud Monitoring: This is a tool for observability of your cloud resources, monitoring performance, and uptime, not for compliance documentation.
D . Security Command Center: This tool provides a centralized view of your security posture, including identifying and managing risks, but it does not offer compliance documentation or reports. Google Cloud Digital Leader Reference:
Refer to Compliance Reports Manager in the Google Cloud documentation for details on accessing compliance and security documentation, including certifications like ISO 27001.

Question: 323

An organization has an ecommerce application that requires a cost-effective, transactional database.

The application will only serve customers in a single region. Which service should they use?

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

Answer: C

Explanation:
The correct answer is C. Cloud SQL. Here's why:
Context of the Questio n : The organization needs a cost-effective, transactional database for an ecommerce application serving customers in a single region. A transactional database is needed to handle online transactions with ACID (Atomicity, Consistency, Isolation, Durability) properties.
Google Cloud Product Relevance:
Cloud SQL is a fully managed relational database service for MySQL, PostgreSQL, and SQL Server. It is cost-effective, supports transactional workloads, and is suitable for applications with regional data residency requirements. This makes it ideal for an e-commerce application that only serves customers in a single region.
Why Not Other Options:
A . Cloud Spanner: Cloud Spanner is a globally distributed, strongly consistent database designed for large-scale, mission-critical applications. It is more expensive and better suited for applications requiring global distribution and high availability across multiple regions.
B . Cloud Bigtable: Cloud Bigtable is a NoSQL wide-column database designed for low-latency, high-throughput workloads, such as analytics. It does not support ACID transactions, making it unsuitable for transactional workloads like an e-commerce application.
D . BigQuery: BigQuery is a fully managed data warehouse designed for large-scale analytics. It is not suitable for transactional database needs.
Google Cloud Digital Leader Reference:
For more details on Cloud SQL, refer to the Cloud SQL documentation in the Google Cloud platform.

Question: 324

An organization wants to adopt the advanced machine learning capabilities of the Google Cloud. However, regulations require data to be stored in an on-premises data center.
Which approach should the organization use?

A. A private-cloud approach
B. A multi-cloud approach
C. A hybrid-cloud approach
D. A public-cloud approach

Answer: C

Explanation:
The correct answer is C. A hybrid-cloud approach. Here's why:
Context of the Questio n : The organization wants to use Google Cloud's advanced machine learning capabilities while maintaining data storage on-premises due to regulatory requirements.
Google Cloud Product Relevance:

A hybrid-cloud approach combines on-premises infrastructure with cloud services. This approach allows organizations to keep sensitive data on-premises while leveraging cloud services for additional computing capabilities, such as advanced machine learning. Google Cloud offers tools like Anthos and Google Cloud's hybrid and multi-cloud solutions to facilitate this integration, enabling the organization to use cloud-based machine learning tools while keeping data in their local data center.

Why Not Other Options:

A . A private-cloud approach: This refers to a cloud environment that is entirely operated within the organization's own infrastructure, which would not provide access to Google Cloud's ML capabilities. B . A multi-cloud approach: While multi-cloud involves using services from multiple cloud providers, it does not specifically address the requirement to keep data on-premises.

D . A public-cloud approach: This would involve moving data to the public cloud, which contradicts the requirement to keep data stored on-premises.

Google Cloud Digital Leader Reference:

Refer to Hybrid and Multi-Cloud Solutions in Google Cloud documentation to understand how Google Cloud can integrate with on-premises data centers while providing cloud services.

Question: 325

What is a defining function of streaming analytics?

A. Processing data records in batches
B. Processing data records continuously
C. Accessing data with high latency
D. Processing a one-off data backfill

Answer: B

Explanation:

Streaming analytics is the real-time processing of data as it is ingested. The defining function of streaming analytics is to process data records continuously, as opposed to batch processing, which handles data in large chunks at intervals. Streaming analytics allows for immediate insights and actions based on live data.

Option B: Processing data records continuously is correct because it describes the fundamental nature of streaming analytics: it analyzes data streams as they come in, enabling real-time analysis, monitoring, and decision-making.

Reference:

Google Cloud Documentation on Dataflow: Real-Time Stream Processing

Google Cloud Platform's Big Data Solutions: Stream Analytics

Question: 326

An organization is building advanced machine learning models in Google Cloud by using TensorFlow. They want to develop their models faster with purpose-built hardware. Which solution should the organization use?

A. CPUs
B. GPUs
C. TPUs
D. DPUs

Answer: C

Explanation:
Tensor Processing Units (TPUs) are a type of purpose-built hardware designed specifically to accelerate machine learning workloads, particularly those involving TensorFlow. TPUs provide significant speed and performance improvements over general-purpose CPUs and even GPUs when running complex ML models.
Option C: TPUs is correct because TPUs (Tensor Processing Units) are optimized to accelerate the training of machine learning models developed with TensorFlow, providing faster performance compared to other hardware.
Reference:
Google Cloud: TPUs Overview
Google Cloud AI and Machine Learning Products: Accelerators for TensorFlow

Question: 327

A retail organization is training a model to recommend products to customers for an ecommerce website. The model was trained on previous purchases, but did not include demographic information on each buyer. What dimension of the data is responsible for the model's poor performance?

A. Accuracy
B. Validity
C. Timeliness
D. Completeness

Answer: D

Explanation:
Completeness refers to the degree to which all necessary data is present and accounted for. In the context of machine learning, if a model is trained without including all relevant dimensions (such as demographic information in this case), it might lead to poor performance because it lacks important data needed to make accurate predictions.
Option D: Completeness is correct because the lack of demographic information represents incomplete data. The absence of this data dimension could prevent the model from learning patterns that are influenced by demographic factors, thereby impacting its overall effectiveness.
Reference:
Google Cloud Documentation: Data Quality in Machine Learning
Google Cloud AI Platform: Best Practices for Training Data Preparation

Question: 328

An organization is concerned about the risk of data loss that may occur due to hardware failures or cyber attacks. They want to restore their systems to a previous state if such an event occurs. What should the organization do?

A. Use Cloud Monitoring.
B. Back up data regularly.
C. Set service level objectives (SLOs).

D. Enable autoscaling.

Answer: B

Explanation:
The correct answer is B. Back up data regularly. Here's why:
Context of the Questio n : The organization is concerned about data loss due to hardware failures or cyber-attacks and wants to restore their systems to a previous state in case of such events.
Google Cloud Product Relevance:
Backing up data regularly is the most effective way to ensure data can be restored in case of loss or damage. Regular backups create copies of data that can be stored separately from the main systems, ensuring that even if the primary data is lost or compromised, the backup can be used to restore it to a previous state. Google Cloud offers several backup solutions, such as Cloud Storage for storing backups and Cloud SQL automated backups for database instances.
Why Not Other Options:
A . Use Cloud Monitoring: Cloud Monitoring helps observe and monitor the health and performance of applications and infrastructure but does not provide backup or data restoration capabilities.
C . Set service level objectives (SLOs): SLOs are targets for service availability and performance but do not directly help with data backup or recovery.
D . Enable autoscaling: Autoscaling helps automatically adjust resources to meet demand but does not protect against data loss or provide restoration capabilities.
Google Cloud Digital Leader Reference:
To understand more about backup strategies in Google Cloud, refer to Backup and DR (Disaster Recovery) documentation.

Question: 329

An organization is deciding on the layout of their resource hierarchy in Google Cloud. They have several projects within a folder. What will happen when user access policies are applied to the folder?

A. The policy applies to the folder only, and will not be inherited by any projects.
B. The policy will be inherited by the projects and their resources within the folder.
C. The policy will be applied to all folders within the organization.
D. The policy will be inherited by the projects in the folder but will not affect their resources.

Answer: B

Explanation:
The correct answer is B. The policy will be inherited by the projects and their resources within the folder.
Here's why:
Context of the Questio n : The organization is deciding on the layout of their resource hierarchy in Google Cloud and wants to understand the impact of applying user access policies at the folder level. Google Cloud Product Relevance:
In Google Cloud, policies applied at a higher level in the resource hierarchy (like an organization or folder) are inherited by all lower-level resources (such as projects and resources within those projects). Therefore, if a policy is applied to a folder, it will be inherited by all projects and their resources under that folder.

This is a core concept in Google Cloud's Identity and Access Management (IAM) system, which ensures that permissions and policies are consistently applied throughout the resource hierarchy. Why Not Other Options:

A . The policy applies to the folder only, and will not be inherited by any projects: This is incorrect because IAM policies are inherited down the resource hierarchy.

C . The policy will be applied to all folders within the organization: This is incorrect because the policy applies only to the specific folder and its contents, not to other folders.

D . The policy will be inherited by the projects in the folder but will not affect their resources: This is incorrect because policies inherited by projects also apply to the resources within those projects.

Google Cloud Digital Leader Reference:

Refer to Google Cloud IAM documentation to understand how IAM policies are inherited in a resource hierarchy.

Question: 330

An organization wants to run their custom application in the cloud in a flexible and scalable way without managing any infrastructure.

Which service model should they use?

A. Infrastructure as a service
B. Platform as a service
C. Network as a service
D. Software as a service

Answer: B

Explanation:

The correct answer is B. Platform as a service. Here's why:

Context of the Questio n : The organization wants to run a custom application in the cloud in a flexible and scalable way without managing any infrastructure.

Google Cloud Product Relevance:

Platform as a Service (PaaS) provides a platform that allows customers to develop, run, and manage applications without dealing with the complexity of building and maintaining the underlying infrastructure.

Google Cloud's App Engine is an example of a PaaS offering. It enables organizations to deploy custom applications in a fully managed environment where scaling, load balancing, and infrastructure management are handled automatically.

Why Not Other Options:

A . Infrastructure as a service: IaaS requires the organization to manage the underlying infrastructure (like virtual machines, networks, and storage), which is not suitable if the organization does not want to manage any infrastructure.

C . Network as a service: NaaS is focused on network connectivity and does not provide the capabilities needed to run custom applications.

D . Software as a service: SaaS provides pre-built applications managed by a provider; it is not suitable for running custom applications.

Google Cloud Digital Leader Reference:

For more on PaaS, refer to the Google App Engine documentation.

Question: 331

An organization is developing a new container-based application. They do not know how popular the application will be when launched and they do not want to pay for idle infrastructure resources.
Which benefit of serverless computing will address this concern?

A. Reduced development costs
B. Built-in security
C. Scalability
D. Disaster recovery

Answer: C

Explanation:

Serverless computing offers automatic scaling, which means it dynamically adjusts the computational resources based on the demand for the application. This benefit is crucial when the popularity of the application is uncertain. The organization only pays for the actual compute resources used, rather than maintaining and paying for idle infrastructure.
Option C: Scalability is correct because serverless computing automatically scales up or down based on demand, ensuring that the organization does not incur costs for idle resources.

Reference:
Google Cloud Serverless Solutions: Benefits of Serverless Computing
Google Cloud Functions and Cloud Run: Auto Scaling Features

Question: 332

An organization needs to store daily transactional data such as customer records and purchase history. The data follows a consistent schema and is cross-referenced. Which type of service should the organization use?

A. Non-relational database
B. Data lake
C. Relational database
D. Data warehouse

Answer: C

Explanation:
Relational databases are ideal for storing transactional data that follows a consistent schema and requires cross-referencing (e.g., customer records and purchase history). They provide ACID (Atomicity, Consistency, Isolation, Durability) compliance, which is essential for reliable transaction processing.
Option C: Relational database is correct because it is specifically designed for structured data with a consistent schema and supports complex queries and cross-referencing of data, which is essential for transactional systems.

Reference:
Google Cloud: Cloud SQL and Cloud Spanner (Relational Database Services)
Google Cloud Database Solutions: Selecting the Right Database

Question: 333

An organization is using new technologies to change its business processes and culture to develop new customer experiences and adapt to market dynamics.
What is the name of this approach?

A. Digital transformation
B. Operations optimization
C. Data center migration
D. User personalization

Answer: A

Explanation:
Digital transformation refers to the use of new technologies to fundamentally change how an organization operates, interacts with customers, and delivers value. It includes changes to business processes, culture, and customer experiences to better adapt to market dynamics and leverage new digital capabilities.
Option A: Digital transformation is correct because it involves using new technologies to drive changes in business processes, customer experiences, and organizational culture to stay competitive in a rapidly changing market.
Reference:
Google Cloud Digital Transformation: Accelerate Business Growth with Cloud
Google Cloud: Digital Transformation Strategy

Question: 334

A real estate organization processes photos of properties uploaded by their agents to an internal application. They must manually check and remove photos that contain personally identifiable information like passports and credit cards. They want a solution that their developer can use to integrate with their app as quickly as possible. Which solution should they use?

A. BigQuery ML
B. Custom training
C. Pre-trainedAPIs
D. AutoML

Answer: C

Explanation:
Pre-trained APIs are ready-made models provided by Google Cloud that can be easily integrated into applications to perform specific tasks, such as detecting and removing personally identifiable information (PII) from images. These APIs, such as Cloud Vision API, can be quickly implemented by developers without needing to build or train custom models.
Option C: Pre-trained APIs is correct because it offers a quick and efficient solution for identifying and removing PII from images, which the organization's developers can easily integrate into their application.
Reference:
Google Cloud Vision API: Pre-trained Models for Image Analysis
Google Cloud AI and Machine Learning APIs: Use Cases and Integration

Question: 335

A cinema company wants to build a model to predict customer visit patterns for the coming year.
They have three years of customer visit data across 300 theaters; however, the data has been stored in different formats by different theaters. They must train the ML model. What should they do?

A. Choose an ML model type that can process different formats of input data.
B. Transform the data into a consistent format.
C. Use the last year of data so there are fewer inconsistencies for the model to handle.
D. Group different format types and train a different model for each group.

Answer: B

Explanation:
The correct answer is B. Transform the data into a consistent format. Here's why:
Context of the Questio n : The cinema company wants to build a machine learning model to predict customer visit patterns using three years of customer visit data stored in different formats. For effective ML model training, the input data must be in a consistent format.
Google Cloud Product Relevance:
To build a robust and accurate ML model, the data used for training needs to be cleaned, preprocessed, and transformed into a consistent format. This process ensures that the model can interpret and learn from the data correctly.
Google Cloud provides services like Dataflow for data transformation and processing, and Dataprep for data cleaning and preparation. These tools help standardize data formats before feeding it into a machine learning model.
Why Not Other Options:
A . Choose an ML model type that can process different formats of input data: Most ML models require data to be in a uniform format; choosing a model that processes multiple formats is not standard practice and can lead to inaccuracies.
C . Use the last year of data so there are fewer inconsistencies for the model to handle: This would limit the amount of data available for training, reducing the model's accuracy and effectiveness.
D . Group different format types and train a different model for each group: Training multiple models for different data formats is inefficient and complex compared to transforming the data into a single consistent format.
Google Cloud Digital Leader Reference:
Refer to Dataflow and Dataprep documentation for details on data transformation and preparation services in Google Cloud.

Question: 336

An organization wants to refactor their application by using a microservices architecture when migrating to the cloud. Which benefit would this action provide?

A. The refactored application is more efficient and scalable.
B. No code changes will be required.
C. This migration pattern provides the fastest path to the cloud.
D. The application will become PCI-DSS compliant by default.

Answer: A

Explanation:

The correct answer is A. The refactored application is more efficient and scalable. Here's why:

Context of the Questio n : The organization wants to refactor its application into a microservices architecture when migrating to the cloud.

Google Cloud Product Relevance:

Microservices architecture breaks down applications into smaller, independent services that can be developed, deployed, and scaled independently. This increases efficiency and scalability, as each microservice can be scaled according to its needs without affecting the entire application.

Google Cloud provides tools such as Kubernetes Engine (GKE) for container orchestration, Anthos for managing microservices across hybrid and multi-cloud environments, and Cloud Run for deploying containerized applications efficiently.

Why Not Other Options:

B . No code changes will be required: Refactoring to microservices often requires significant code changes to decouple services.

C . This migration pattern provides the fastest path to the cloud: Refactoring to microservices is a complex process that typically takes more time compared to a "lift-and-shift" migration.

D . The application will become PCI-DSS compliant by default: Compliance with PCI-DSS (Payment Card Industry Data Security Standard) depends on specific security controls, not on using microservices architecture.

Google Cloud Digital Leader Reference:

For more information on microservices and cloud-native architectures, refer to the Google Cloud Microservices documentation.

Question: 337

An organization supplies electric car chargers in a nationwide network. They store customer and charging data in BigQuery. They want to efficiently upgrade and maintain their network to reduce waste and achieve their sustainability goals. How can Google Cloud services help the organization?

A. Secure data access by using Identity and Access Management.
B. Move data from BigQuery to Cloud Storage to reduce storage costs.
C. Provide access to third-party developers by creating an API using Apigee.
D. Create a model in BiaQuery to Dredict future maintenance schedules.

Answer: D

Explanation:

The correct answer is D. Create a model in BigQuery to predict future maintenance schedules. Here's why:

Context of the Questio n : The organization wants to use Google Cloud services to efficiently upgrade and maintain their network to reduce waste and achieve sustainability goals.

Google Cloud Product Relevance:

BigQuery is a fully managed data warehouse that supports SQL queries for analyzing large datasets.

It also has built-in machine learning capabilities (BigQuery ML) that allow users to create and train models directly within the database using SQL.

By creating a predictive model in BigQuery, the organization can analyze historical data to forecast future maintenance needs, optimize upgrade schedules, and minimize waste, aligning with their sustainability

goals.

Why Not Other Options:

A . Secure data access by using Identity and Access Management: While IAM is important for securing access, it does not directly help with maintenance prediction or sustainability goals.

B . Move data from BigQuery to Cloud Storage to reduce storage costs: This option focuses on cost savings, not on improving operational efficiency or achieving sustainability goals.

C . Provide access to third-party developers by creating an API using Apigee: This is not directly related to the goal of optimizing maintenance and reducing waste.

Google Cloud Digital Leader Reference:

Refer to BigQuery ML documentation to learn how to create and use machine learning models within BigQuery.

Question: 338

An organization runs a batch data analysis workload on a virtual machine (VM). The workload can be easily restarted without losing work, and is not time critical. Organizations must choose the lowest cost option to run the workload. What option should they choose?

A. A standard VM in a pay-as-you-go model on Compute Engine
B. A Cloud Function with a small memory limit
C. A Preemptible or Spot VM on Compute Engine
D. A custom VM in a pay-as-you-go model on Compute Engine

Answer: C

Explanation:

The correct answer is C. A Preemptible or Spot VM on Compute Engine. Here's why:

Context of the Questio n : The organization runs a batch data analysis workload that is not time- critical and can be restarted without losing work, and they want to choose the lowest-cost option. Google Cloud Product Relevance:

Preemptible VMs (also known as Spot VMs) are short-lived compute instances that offer a significantly lower cost than standard VMs. They are ideal for non-time-critical workloads that can be interrupted and restarted, such as batch data analysis.

Spot VMs are a cost-effective choice for workloads that can tolerate interruptions, providing substantial cost savings compared to standard VMs.

Why Not Other Options:

A . A standard VM in a pay-as-you-go model on Compute Engine: This would be more expensive than using preemptible or spot VMs.

B . A Cloud Function with a small memory limit: Cloud Functions are designed for event-driven tasks, not for long-running batch data analysis.

D . A custom VM in a pay-as-you-go model on Compute Engine: This would also be more expensive compared to using preemptible or spot VMs.

Google Cloud Digital Leader Reference:

Refer to Compute Engine Spot VMs documentation to understand how to use and benefit from Spot VMs for cost savings.

Question: 339

What is the purpose of an application programming interface (API)?

A. To provide cloud plugins for integrated development environments
B. To manage multiple containerized workloads
C. To connect networks from different cloud providers
D. To provide a set of instructions that allow computer programs to communicate with each other

Answer: D

Explanation:
An Application Programming Interface (API) is a set of protocols, routines, and tools that allow different software applications to communicate with each other. APIs provide the necessary instructions for one program to interact with another, enabling the integration of different systems and the development of new functionalities.
Option D: To provide a set of instructions that allow computer programs to communicate with each other is correct because it accurately describes the fundamental purpose of an API.
Reference:
Google Cloud: Understanding APIs and their Use Cases
Google Cloud API Documentation: API Basics

Question: 340

When customer data is uploaded to Google Cloud, who owns the data?

A. The customer and Google share ownership
B. The customer
C. Google
D. A third party

Answer: B

Explanation:
When data is uploaded to Google Cloud, the customer retains full ownership of their data. Google Cloud acts as the data processor, providing the infrastructure and services for storage, but the rights to the data, including access, control, and management, remain with the customer.
Option B: The customer is correct because, according to Google Cloud's data privacy policy, the customer owns their data.
Reference:
Google Cloud Privacy Policy: Data Ownership and Responsibility
Google Cloud Trust Principles: Data Protection and Privacy

Question: 341

An organization is deploying their servers to the cloud using the infrastructure as a service model.
In the shared responsibility model, what is the cloud provider responsible for?

A. Security of the operating system
B. Physical security

C. Data access policies
D. Security of the software

Answer: B

Explanation:
In the Infrastructure as a Service (IaaS) model under the shared responsibility model, the cloud provider (e.g., Google Cloud) is responsible for the physical security of the infrastructure, including data centers, hardware, and the physical environment. The customer is responsible for managing the operating system, applications, and data.
Option B: Physical security is correct because it describes the cloud provider's responsibility in securing the physical aspects of the infrastructure.
Reference:
Google Cloud Shared Responsibility Model: Customer and Provider Responsibilities
Google Cloud Security Overview: Physical Security Measures

Question: 342

An organization is using Compute Engine and wants to receive sustained-use discounts. What should the organization do?

A. Choose preemptible or spot instances when creating virtual machines.
B. Nothing. Sustained use discounts are automatically applied.
C. Commit to virtual machine usage for a one or three-year period.
D. Choose low-powered virtual machines.

Answer: B

Explanation:
Sustained-use discounts are automatically applied by Google Cloud when Compute Engine instances are run for a significant portion of the month. These discounts do not require any prior commitment or action from the customer; they are automatically calculated based on usage.
Option B: Nothing. Sustained use discounts are automatically applied is correct because Google Cloud automatically grants these discounts based on usage patterns.
Reference:
Google Cloud Pricing: Sustained Use Discounts Explanation
Google Cloud Compute Engine Pricing Guide

Question: 343

An organization is running critical workloads in production and requires a Google Cloud support service with fast response times and a dedicated Technical Account Manager. Which customer care service level should the organization choose?

A. Premium
B. Standard
C. Enhanced

D. Basic

Answer: A

Explanation:
The Premium customer care service level in Google Cloud provides the fastest response times and includes a dedicated Technical Account Manager (TAM) to assist with managing critical workloads and ensuring smooth operations.
Option A: Premium is correct because it provides the required support features, including fast response times and a dedicated TAM.
Reference:
Google Cloud Support Services: Premium Support Overview
Google Cloud Customer Care Options: Comparison of Support Levels

Question: 344

An organization has a large archive of unstructured data, including video and audio files. Which storage solution should the organization use?

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

Answer: D

Explanation:
The correct answer is D. Cloud Storage. Here's why:
Context of the Questio n : The organization has a large archive of unstructured data, including video and audio files. These types of files are typically large and require scalable storage that supports unstructured data formats.
Google Cloud Product Relevance:
Cloud Storage is a fully managed, scalable, and highly durable object storage service that supports storing any amount of unstructured data, including images, videos, audio files, backups, and more. It offers various storage classes to optimize costs based on access frequency, making it ideal for large archives.
Cloud Storage provides features like versioning, lifecycle management, and access control, which are essential for managing a large archive of files.
Why Not Other Options:
A . Cloud SQL: Cloud SQL is a managed relational database service that supports structured data with predefined schemas, not suitable for storing unstructured data like videos and audio files.
B . Cloud Spanner: Cloud Spanner is a fully managed, globally distributed relational database. It is also designed for structured data and not suitable for large-scale unstructured data storage.
C . Cloud Bigtable: Cloud Bigtable is a NoSQL database for large-scale analytics but is not optimized for storing unstructured data like multimedia files.
Google Cloud Digital Leader Reference:
Refer to Google Cloud Storage documentation for details on storing unstructured data and managing large archives.

Question: 345

What is the Site Reliability Engineering (SRE) term for an organizations desired level of reliability and performance?

A. Enhanced support
B. Scalable infrastructure
C. Service-level indicator
D. Service-level objective

Answer: D

Explanation:
The correct answer is D. Service-level objective. Here's why:
Context of the Questio n : The organization wants to understand the term used in Site Reliability Engineering (SRE) for defining the desired level of reliability and performance.
Google Cloud Product Relevance:
A Service-Level Objective (SLO) is a key term in SRE, representing the target level of reliability or performance for a specific service. SLOs define acceptable levels of service in terms of availability, latency, or other performance metrics. They are used to set clear expectations and help measure whether services meet the desired reliability levels.
SLOs are closely tied to Service-Level Agreements (SLAs) and Service-Level Indicators (SLIs), where SLIs measure specific aspects of service performance, and SLOs set the target levels based on those indicators.
Why Not Other Options:
A . Enhanced support: This is a support option for Google Cloud customers and is not related to SRE concepts.
B . Scalable infrastructure: This refers to cloud infrastructure's ability to scale resources up or down but is not specific to reliability and performance metrics.
C . Service-level indicator: An SLI measures a specific aspect of the service's performance but does not define the target level of performance.
Google Cloud Digital Leader Reference:
Refer to the Site Reliability Engineering (SRE) documentation to learn more about SLOs, SLIs, and SLAs.

Question: 346

An organization is operating multiple workloads in containers and requires full control of how the workloads are configured. Which Google Cloud service should the organization use?

A. Cloud Functions
B. Compute Engine
C. Kubernetes Engine
D. Cloud Run

Answer: C

Explanation:
The correct answer is C. Kubernetes Engine. Here's why:
Context of the Questio n : The organization needs to operate multiple workloads in containers with full control over how these workloads are configured.

Google Cloud Product Relevance:

Google Kubernetes Engine (GKE) is a fully managed service that allows organizations to run, deploy, and manage containerized applications using Kubernetes. GKE provides full control over workload configuration, including resource allocation, scaling, networking, and security policies.

With GKE, the organization can use Kubernetes features such as custom resource definitions, namespaces, and pod configurations to achieve the desired level of control over their workloads.

Why Not Other Options:

A . Cloud Functions: Cloud Functions is a serverless compute service for running event-driven functions and does not provide control over containerized workloads.

B . Compute Engine: Compute Engine provides virtual machines (VMs), which are not specific to containerized workloads. It requires more management compared to GKE.

D . Cloud Run: Cloud Run is a managed service for running containers but abstracts much of the configuration details to simplify deployment. It does not offer as much control over workload configuration as GKE.

Google Cloud Digital Leader Reference:

Refer to the Google Kubernetes Engine documentation for more details on managing containerized workloads.

Question: 347

An organization is developing applications by using Kubernetes. They want their teams to spend more time developing rather than managing clusters. Which Google Cloud service should the organization choose?

A. Cloud Run
B. Compute Engine
C. GKE Autopilot
D. GKE Standard

Answer: C

Explanation:

The correct answer is C. GKE Autopilot. Here's why:

Context of the Questio n : The organization is using Kubernetes and wants to minimize cluster management so their teams can focus more on developing applications.

Google Cloud Product Relevance:

GKE Autopilot is a fully managed mode of operation for Google Kubernetes Engine (GKE) that abstracts much of the cluster management, such as provisioning and managing the control plane, nodes, and infrastructure. It automates operational tasks and optimizes resource management, enabling developers to focus on writing and deploying code rather than managing Kubernetes clusters.

GKE Autopilot is ideal for organizations that want the benefits of Kubernetes without the overhead of managing cluster infrastructure.

Why Not Other Options:

A . Cloud Run: While Cloud Run is a fully managed service for running containers, it is not specifically a Kubernetes service and does not offer the same Kubernetes-native experience as GKE Autopilot.

B . Compute Engine: Compute Engine requires manual management of VMs and is not specific to containerized or Kubernetes applications.

D . GKE Standard: GKE Standard provides more control over cluster configuration but requires more management effort compared to GKE Autopilot.

Google Cloud Digital Leader Reference:

For more details on GKE Autopilot, refer to the GKE Autopilot documentation.

Question: 348

An organization has collected petabytes of historical dat

a. They need an advanced analysis solution that is fast, scalable, and fully managed. Which Google product or service should the organization use?

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

Answer: A

Explanation:

BigQuery is a fully managed, serverless data warehouse that is designed for large-scale data analysis. It is optimized for handling petabytes of data and provides fast query performance. It is scalable and supports advanced analytical queries, making it ideal for organizations that need to analyze massive amounts of historical data efficiently.

Option A: BigQuery is correct because it is specifically designed for large-scale, advanced data analytics and is fully managed by Google Cloud.

Reference:

Google Cloud: BigQuery Product Overview
Google Cloud BigQuery: Advanced Data Analysis Capabilities

Question: 349

An organization is planning to deploy a new workload to Google Cloud. They need an accurate estimate of the likely costs of running the workload. How should the organization create this estimate?

A. Use the Google Cloud Pricing Calculator.
B. Deploy workload to test environment to observe costs.
C. Use historic costs of an existing similar workload.
D. Refer to pricing information and manually calculate an estimate.

Answer: A

Explanation:

The Google Cloud Pricing Calculator is a tool designed to help users estimate the cost of Google Cloud services based on their specific needs and configurations. It allows organizations to input various service usage parameters, such as storage, compute, and networking requirements, to calculate a detailed cost estimate.

Option A: Use the Google Cloud Pricing Calculator is correct because it provides an accurate and comprehensive way to estimate the costs of running a workload on Google Cloud.

Reference:

Google Cloud Pricing Calculator: Cost Estimation Tool
Google Cloud Cost Management: How to Use the Pricing Calculator

Question: 350

An organization is evaluating its defenses against cyber security threats and is concerned about the risks of social engineering by cyber criminals. How might these attacks happen?

A. Phishing emails
B. SQL injection attacks
C. Physical damage to hardware
D. Distributed denial-of-service attacks

Answer: A

Explanation:
Phishing emails are a common form of social engineering attack where attackers impersonate legitimate entities to trick users into revealing sensitive information, such as passwords or credit card details. These attacks exploit human psychology rather than technical vulnerabilities.
Option A: Phishing emails is correct because it represents a primary method of social engineering used by cybercriminals to manipulate individuals into divulging confidential information.
Reference:
Google Cloud Security Best Practices: Social Engineering Attacks
Google Cloud Cybersecurity Framework: Protection Against Phishing

Question: 351

An organization needs protection against distributed denial-of-service (DDoS) attacks. Which Google Cloud service should the organization use?

A. Cloud Build
B. Google Cloud Armor
C. Security Command Center
D. Cloud VPN

Answer: B

Explanation:
Google Cloud Armor is a service that provides protection against Distributed Denial-of-Service (DDoS) attacks. It offers a scalable web application firewall (WAF) that helps to defend against layer 3,

4, and 7 DDoS attacks, ensuring that applications remain available and responsive.
Option B: Google Cloud Armor is correct because it is specifically designed to provide protection against DDoS attacks.
Reference:
Google Cloud Armor: DDoS Protection and WAF Features
Google Cloud Security: Overview of DDoS Protection Services

Question: 352

An organization is deploying applications to the cloud using platform as a service.

In the shared responsibility model, what remains the responsibility of the customer organization?

A. Physical security
B. Network security
C. Operations
D. Data access policies

Answer: D

Explanation:
The correct answer is D. Data access policies. Here's why:
Context of the Questio n : The organization is deploying applications to the cloud using Platform as a Service (PaaS), which means certain responsibilities are shared between the cloud provider (Google Cloud) and the customer.
Google Cloud Product Relevance:
In the shared responsibility model for PaaS, Google Cloud manages the underlying infrastructure, including physical security, network security, and operations. However, the customer organization is still responsible for managing their data, including defining and enforcing data access policies. Data access policies involve setting permissions, defining roles, and controlling who can access specific data within the organization, which is a critical component of maintaining data security and privacy.
Why Not Other Options:
A . Physical security: This is managed by the cloud provider (Google Cloud) and includes securing data centers and physical hardware.
B . Network security: While the cloud provider ensures network security within its infrastructure, the customer may have some responsibilities at the application level, but in a PaaS environment, most network security responsibilities are managed by the provider.
C . Operations: In a PaaS model, many operational aspects such as patch management, updates, and scaling are handled by the cloud provider.
Google Cloud Digital Leader Reference:
Refer to Google Cloud Shared Responsibility Model documentation for a deeper understanding of responsibilities between the provider and the customer.

Question: 353

When is data automatically encrypted in Google Cloud?
A. When it is at rest only.
B. When it is at rest and in transit.
C. When it is in transit only.
D. Data is not automatically encrypted by default.

Answer: B

Explanation:
The correct answer is B. When it is at rest and in transit. Here's why:
Context of the Questio n : The question is about when data is automatically encrypted by Google Cloud.
Google Cloud Product Relevance:
Google Cloud automatically encrypts data at rest and in transit. This means that data stored in Google Cloud (such as in Google Cloud Storage or BigQuery) is encrypted by default to protect it from unauthorized access.

Similarly, data in transit between Google Cloud services or between users and Google Cloud is encrypted using TLS (Transport Layer Security).
These automatic encryption measures are part of Google's comprehensive approach to data security, ensuring that data remains protected throughout its lifecycle.
Why Not Other Options:
A . When it is at rest only: This is incorrect because data is also encrypted in transit.
C . When it is in transit only: This is incorrect because data is also encrypted at rest.
D . Data is not automatically encrypted by default: This is incorrect as Google Cloud provides automatic encryption for data both at rest and in transit.
Google Cloud Digital Leader Reference:
Refer to Google Cloud Encryption documentation for more information on how data is encrypted in Google Cloud.

Question: 354

A vacation home rental organization wants to predict the popularity of properties in their upcoming busy season. They do not have a data science team, and want to use their in-house database administration skills to create a machine learning model. What should the organization do?

A. Use custom training in Vertex Al with TensorFlow.
B. Use BigQuery ML and create models using SQL.
C. Build a model in AutoML using labeled data.
D. Integrate pre-trained APIs into their application.

Answer: B

Explanation:

The correct answer is B. Use BigQuery ML and create models using SQL. Here's why:
Context of the Questio n : The organization wants to predict property popularity for the upcoming busy season but does not have a data science team and relies on in-house database administration skills.
Google Cloud Product Relevance:
BigQuery ML enables data analysts and database administrators to build and run machine learning models using SQL queries directly within BigQuery. This makes it accessible to users familiar with SQL, eliminating the need for advanced data science skills.
BigQuery ML is ideal for this organization since it allows them to leverage their existing SQL skills to build predictive models without the complexity of custom machine learning frameworks or coding in languages like Python.
Why Not Other Options:
A . Use custom training in Vertex AI with TensorFlow: This requires expertise in TensorFlow and machine learning, which the organization lacks.
C . Build a model in AutoML using labeled data: AutoML does simplify the ML model-building process but still requires more knowledge than basic SQL. Also, BigQuery ML allows the use of existing SQL skills, which is preferred in this scenario.
D . Integrate pre-trained APIs into their application: Pre-trained APIs are suitable for common use cases like image recognition or natural language processing, not for custom predictions based on unique data like property popularity.
Google Cloud Digital Leader Reference:

Refer to BigQuery ML documentation for more details on using SQL to create and use ML models.

Question: 355

An organization wants to ensure that they grant users only the permissions they require to perform their job responsibilities. Which security principle describes this approach?

A. Security by default
B. Least privilege
C. Zero-trust
D. Cyber resilience

Answer: B

Explanation:
The correct answer is B. Least privilege. Here's why:
Context of the Questio n : The organization wants to ensure that users are granted only the permissions required for their job responsibilities, aligning with a specific security principle. Google Cloud Product Relevance:
The Least Privilege principle is a fundamental security concept that dictates that users should be given the minimum level of access—or permissions—necessary to perform their job functions. This reduces the risk of unauthorized access and potential security breaches.
In Google Cloud, Identity and Access Management (IAM) can be used to implement the least privilege principle by assigning roles and permissions that are tightly scoped to users' specific responsibilities.
Why Not Other Options:
A . Security by default: This is a general security approach where security settings are enabled by default, but it does not specifically address access control.
C . Zero-trust: Zero-trust is a broader security model that assumes no implicit trust within a network; it complements but does not replace the least privilege approach.
D . Cyber resilience: This refers to an organization’s ability to continue operations during a cyber incident, which is different from granting minimal permissions.
Google Cloud Digital Leader Reference:
Refer to Google Cloud IAM documentation for more information on implementing the principle of least privilege in Google Cloud.

Question: 356

An organization wants to analyze data in a data warehouse. How should they proceed?

A. Choose a system to store structured and semi-structured data that supports ad-hoc analysis and custom reporting.
B. Copy unstructured data into a single large object store.
C. Ensure data is stored in structured tables and rows to support transactional queries and relationships.
D. Import data into a semi-structured time-series database.

Answer: A

Explanation:

To analyze data in a data warehouse, the organization should use a system that can handle both structured and semi-structured data while supporting ad-hoc analysis and custom reporting. A data warehouse like BigQuery is designed to efficiently store and query large amounts of structured and semi-structured data, allowing for flexible and real-time analytical queries.

Option A: Choose a system to store structured and semi-structured data that supports ad-hoc analysis and custom reporting is correct because it aligns with the capabilities needed for effective data warehousing and analysis.

Reference:

Google Cloud BigQuery: Data Warehousing and Analysis

Google Cloud Data Solutions: Ad-hoc Analysis and Reporting Tools

Question: 357

An organization wants to build a data pipeline to transform its data so it can be reconciled in a data warehouse. The solution must be scalable and require little or no management. Which Google product or service should the organization choose?

A. Pub/Sub
B. Dataflow
C. Cloud Storage
D. Cloud Bigtable

Answer: B

Explanation:

Google Cloud Dataflow is a fully managed service for creating data pipelines that can handle both batch and stream processing. It is scalable and requires minimal management, making it ideal for transforming data before reconciling it in a data warehouse. Dataflow integrates seamlessly with other Google Cloud services, such as BigQuery, to create end-to-end data processing solutions.

Option B: Dataflow is correct because it provides a scalable and managed solution for building data pipelines.

Reference:

Google Cloud Dataflow: Scalable Data Processing Service

Google Cloud Data Integration: Building Data Pipelines with Dataflow

Question: 358

structured data?

A. Customer orders
B. Product ratings
C. Historical stock inventory
D. Call center transcripts

Answer: D

Explanation:
Unstructured data refers to data that does not have a pre-defined data model or schema. Examples include text, images, audio, and video files. Call center transcripts are unstructured data because they are composed of free-form text, which does not follow a specific schema or structure.
Option D: Call center transcripts is correct because it represents unstructured data, as opposed to structured data like customer orders or product ratings.
Reference:
Google Cloud: Understanding Unstructured Data
Google Cloud Storage Solutions: Managing Unstructured Data

Question: 359

A financial organization has many customers who close their accounts every year. The organization wants to use data and AI to identify at-risk customers, so they can retain customers by offering discounts and improved services. What should the organization do?

A. Create a survey for all customers to identify their current level of satisfaction.
B. Create a report based on last year's customer feedback.
C. Create a dashboard of previous customers that have exited, and look for obvious correlations in the visualization.
D. Create a ML model based on the demographics and activities of previous customers that exited.

Answer: D

Explanation:
Creating a machine learning (ML) model to analyze the demographics and behaviors of customers who have previously closed their accounts is an effective approach to identify at-risk customers. This model can predict which current customers are most likely to leave, allowing the organization to proactively offer discounts or improved services to retain them.
Option D: Create a ML model based on the demographics and activities of previous customers that exited is correct because it uses data-driven insights and AI to identify patterns and predict future customer behavior.
Reference:
Google Cloud AI Platform: Building Machine Learning Models
Google Cloud Customer Retention Solutions: Using AI and Machine Learning

Question: 360

A manufacturing organization has a large collection of images labeled as intact or defective parts. They want to use this data to build a simple solution to detect faulty parts on their production line. They have no data science expertise. Which solution should they use?

A. Discovery AI for Retail
B. Pre-trainedAPIs
C. Document AI
D. AutoML

Answer: D

Explanation:

The correct answer is D. AutoML. Here's why:
Context of the Questio n : The organization wants to build a solution to detect faulty parts on a production line using a large collection of labeled images (intact or defective parts). They do not have data science expertise, so they need a tool that simplifies the machine learning process.
Google Cloud Product Relevance:
AutoML is a suite of machine learning products that allows users to build custom models easily without needing deep expertise in machine learning or data science. It is designed to simplify the process of training, evaluating, and deploying models, especially in cases like image recognition where labeled datasets are already available.
AutoML Vision, a part of AutoML, would be ideal for this use case as it specifically handles image classification and can easily differentiate between intact and defective parts.
Why Not Other Options:
A . Discovery AI for Retail: This is a solution tailored for retail use cases like product discovery and search optimization, not for manufacturing defect detection.
B . Pre-trained APIs: While pre-trained APIs (like Vision API) can recognize general image patterns, they may not be specific enough for a custom use case like detecting defective parts.
C . Document AI: This service is designed for understanding and processing documents, not for analyzing images.
Google Cloud Digital Leader Reference:
For more information on AutoML, refer to the AutoML Vision documentation.

Question: 361

Customers are reporting very high latencies when accessing an application from the United States. The application is currently running in a single region in Europe.
What should the organization do?
A. Set up a new billing account in the United States.
B. Run the application in additional regions in Europe.
C. Run the application in additional zones in the European region.
D. Run a replica of the application in a region in the United States.

Answer: D

Explanation:
The correct answer is D. Run a replica of the application in a region in the United States. Here's why: Context of the Questio n : The application currently runs in a single region in Europe, and customers in the United States are experiencing high latencies. The organization needs to reduce latencies for these customers.
Google Cloud Product Relevance:
Running a replica of the application in a region in the United States would reduce latency by hosting the application closer to the end users. This improves response times and ensures a better user experience for customers in the United States.
Google Cloud provides services such as Cloud Load Balancing to distribute traffic efficiently between instances in different regions.
Why Not Other Options:

A . Set up a new billing account in the United States: Billing accounts do not affect application performance or latency.

B . Run the application in additional regions in Europe: This would not reduce latencies for users in the United States, as their traffic would still need to travel across the Atlantic.

C . Run the application in additional zones in the European region: Additional zones in Europe would not help with reducing latency for U.S. customers.

Google Cloud Digital Leader Reference:

Refer to Google Cloud's multi-region deployment documentation for more information on deploying applications across multiple regions.

Question: 362

An organization processes batch sales data at the end of every month to analyze sales trends and derive business insights. They want to improve accuracy and make near real-time decisions. What should the organization do?

A. Filter the data so reports are generated faster.
B. Switch from batch processing to stream processing.
C. Process batch reports weekly instead of monthly.
D. Change from a relational database to a NoSQL database.

Answer: B

Explanation:

The correct answer is B. Switch from batch processing to stream processing. Here's why:

Context of the Questio n : The organization processes sales data at the end of every month and wants to make near real-time decisions to improve accuracy and agility.

Google Cloud Product Relevance:

Stream processing allows data to be processed as it arrives, enabling real-time analytics and faster decision-making. This would allow the organization to analyze sales trends continuously rather than waiting for monthly batch jobs to complete.

Google Cloud services like Dataflow support both batch and stream processing, allowing for the transformation and analysis of data in real-time.

Why Not Other Options:

A . Filter the data so reports are generated faster: Filtering data may speed up reporting slightly but does not enable near real-time decision-making.

C . Process batch reports weekly instead of monthly: Processing weekly instead of monthly still doesn't provide the near real-time insights the organization needs.

D . Change from a relational database to a NoSQL database: This change might improve scalability or flexibility but does not directly address the need for real-time data processing.

Google Cloud Digital Leader Reference:

Refer to Google Cloud Dataflow documentation to learn about stream processing capabilities.

Question: 363

An organization wants to control what types of network traffic are allowed to enter and leave its network and access its applications. Which security measure should the organization use?

A. Encryption keys
B. 2-Step Verification
C. Firewall rules
D. Privileged access

Answer: C

Explanation:
The correct answer is C. Firewall rules. Here's why:

Context of the Questio n : The organization wants to control what types of network traffic are allowed to enter and leave its network and access its applications.

Google Cloud Product Relevance:

Firewall rules are security settings used to control the flow of traffic to and from instances or resources in a network. They define which traffic is allowed or denied based on parameters such as IP address, protocol, and port.

Firewall rules are essential for protecting applications from unauthorized access and for ensuring that only trusted traffic is permitted.

Why Not Other Options:

A . Encryption keys: Encryption keys are used for data protection and encryption but do not control network traffic.

B . 2-Step Verification: This enhances account security by requiring a second form of authentication but does not control network traffic.

D . Privileged access: This refers to managing access rights for administrative tasks and is not related to controlling network traffic.

Google Cloud Digital Leader Reference:

Refer to Google Cloud VPC firewall documentation to understand how firewall rules are used to control network traffic.

Question: 364

An organization has a large dataset that contains text transcripts of conversations between their customers and service representatives. They want an automated solution to identify the topics their customers care most about. Which service should the organization use?

A. Cloud Translation API
B. Speech-to-Text API
C. Vision API
D. Natural Language API

Answer: D

Explanation:
The Natural Language API from Google Cloud is designed to understand and analyze human language. It can automatically identify and extract key topics, sentiments, and entities from text data, making it ideal for processing large datasets of text transcripts to determine what topics customers care most about.

Option D: Natural Language API is correct because it is specifically built for text analysis, topic modeling, and sentiment analysis.

Reference:
Google Cloud Natural Language API: Text Analysis and Entity Recognition
Google Cloud AI Solutions: Using Natural Language Processing for Customer Insights

Question: 365

An organization is hosting an application in Europe, and customers in Asia are reporting slow response times despite their fast internet connection.
What is the problem?

A. Misconfigured application servers
B. Network latency
C. Not enough application servers
D. Network bandwidth

Answer: B

Explanation:
Network latency refers to the delay that occurs in data transmission over a network, which can affect the response time of an application for users who are geographically distant from the server hosting the application. In this case, customers in Asia experience slow response times because the application is hosted in Europe, causing a higher latency due to the physical distance between the user and the server.
Option B: Network latency is correct because it addresses the primary issue of increased delay due to geographical distance.
Reference:
Google Cloud Networking: Understanding Network Latency and Performance

Google Cloud: Global Load Balancing Solutions to Reduce Latency

Question: 366

An organization has a large VMWare environment that they want to migrate to the cloud. They want to retain existing operational processes and tools. Which Google Cloud service should the organization use?

A. Google Cloud VMware Engine
B. Bare Metal Solution
C. VMWare vSphere
D. Compute Engine

Answer: A

Explanation:
Google Cloud VMware Engine is a fully managed service that allows organizations to run VMware environments natively in Google Cloud. It enables customers to retain their existing VMware operational processes, tools, and policies while migrating to the cloud, providing a seamless transition for VMware-based workloads.
Option A: Google Cloud VMware Engine is correct because it supports VMware environments and allows for continuity of existing operations and tools.
Reference:

Google Cloud VMware Engine: Migrate and Run VMware Workloads Natively
Google Cloud: Hybrid Cloud Solutions for VMware

Question: 367

What is a benefit of Google's purpose-built servers compared to standard servers?

A. They run software that cannot be deployed on standard servers.
B. They are backward compatible with legacy disk drives.
C. They are optimized for specific tasks making them more efficient.
D. They are cheaper to build than standard servers.

Answer: C

Explanation:
Google's purpose-built servers are specifically designed and optimized for Google's workloads, making them more efficient than standard off-the-shelf servers. These servers are tailored for specific tasks, such as handling large-scale data processing or machine learning workloads, providing better performance and energy efficiency.
Option C: They are optimized for specific tasks making them more efficient is correct because it describes the primary benefit of Google's custom server hardware.
Reference:
Google Cloud Data Center Innovation: Purpose-built Servers
Google Cloud Infrastructure: Hardware and Software Optimization

Question: 368

An organization has migrated several large databases to the cloud. It wants to increase the value of its data, improve cost controls and strengthen regulatory compliance. What should the organization do?

A. Delete data stored for over a year.
B. Establish an effective data governance program.
C. Export relational data to modern NoSQL databases.
D. Create monthly reports on data access and uses.

Answer: B

Explanation:
An effective data governance program helps organizations manage data quality, privacy, and security, ensuring compliance with regulations and optimizing the value derived from data. This approach involves defining policies and procedures for data access, usage, and retention, which improves cost control and strengthens regulatory compliance.
Option B: Establish an effective data governance program is correct because it provides a comprehensive framework for managing data, optimizing value, controlling costs, and ensuring compliance.
Reference:
Google Cloud: Data Governance and Compliance Solutions
Google Cloud Data Management: Best Practices for Data Governance

Question: 369

An organization is concerned about their cloud costs. They want to be informed when their spending exceeds a specific threshold, rather than waiting to see their bill at the end of the month. What should the organization do?

A. Configure a budget threshold rule and alert.
B. Pause virtual machines during non-business hours.
C. Adjust project resource quota policies.
D. Use historical cost data to predict future overspend.

Answer: A

Explanation:
The correct answer is A. Configure a budget threshold rule and alert. Here's why:
Context of the Questio n : The organization wants to be informed when their cloud spending exceeds a specific threshold to manage costs proactively rather than finding out at the end of the month.
Google Cloud Product Relevance:
Google Cloud offers a budget and alerting feature in the Cloud Billing platform. Organizations can set up a budget for their cloud expenditures and create threshold alerts that notify them when spending reaches a specified percentage of the budget (e.g., 50%, 90%, or 100%).
This feature helps organizations monitor their cloud costs in real-time and take action to prevent overspending, ensuring better financial management and cost optimization.
Why Not Other Options:
B . Pause virtual machines during non-business hours: This action could save costs but does not provide an alert mechanism based on budget thresholds.
C . Adjust project resource quota policies: Adjusting resource quotas can limit resource use but does not provide cost monitoring or alerting capabilities.
D . Use historical cost data to predict future overspend: While historical data can help with budgeting and forecasting, it does not provide real-time alerts for current spending.
Google Cloud Digital Leader Reference:
Refer to Google Cloud's Budget and alerts documentation for more details on configuring budget alerts.

Question: 370

What is generative AI?

A. A specific field of Al that can be used to create content
B. An umbrella term that includes all types of machine learning algorithms
C. An Al model for labeling images
D. A machine learning model for language translation

Answer: A

Explanation:
The correct answer is A. A specific field of AI that can be used to create content. Here's why:
Context of the Questio n : The question is asking about generative AI, a specialized area within the broader field of artificial intelligence.

Google Cloud Product Relevance:

Generative AI refers to a branch of artificial intelligence that focuses on generating new content, such as text, images, audio, and other data. It uses machine learning models, such as generative adversarial networks (GANs) and transformer models (like GPT), to create new data that mimics the patterns of existing data. This type of AI is used in various applications, including text generation, content creation, image synthesis, and more.

Why Not Other Options:

B . An umbrella term that includes all types of machine learning algorithms: This is incorrect because generative AI is a specific field within AI, not an umbrella term for all AI types.

C . An AI model for labeling images: Generative AI creates content, while labeling images falls under supervised learning.

D . A machine learning model for language translation: Language translation is a different AI application and does not specifically involve generating new content.

Google Cloud Digital Leader Reference:

Refer to the Google Cloud Generative AI documentation for more information on generative AI models and applications.

Question: 371

What is a benefit of the OpEx model for cloud security?

A. Organizations can deploy custom security hardware.
B. Organizations do not need to make upfront capital investments in cloud security.
C. Organizations do not need to configure any security settings for cloud resources.
D. The cloud provider guarantees security.

Answer: B

Explanation:

The correct answer is B. Organizations do not need to make upfront capital investments in cloud security. Here's why:

Context of the Questio n : The question asks about the benefits of the operational expenditure (OpEx) model for cloud security.

Google Cloud Product Relevance:

The OpEx model involves paying for cloud resources and services on a pay-as-you-go basis, avoiding the need for large upfront capital expenditures (CapEx) for security hardware or software.

In the cloud, security services (such as firewalls, encryption, and monitoring tools) are provided as managed services, which organizations can subscribe to and pay for based on usage. This model allows for more flexibility and scalability in managing security needs and costs.

Why Not Other Options:

A . Organizations can deploy custom security hardware: In the cloud, most security needs are met through managed services rather than deploying custom hardware.

C . Organizations do not need to configure any security settings for cloud resources: This is incorrect, as organizations are responsible for configuring security settings for their applications and data.

D . The cloud provider guarantees security: The cloud provider guarantees the security of the infrastructure, but the customer is responsible for securing their applications, data, and access controls.

Google Cloud Digital Leader Reference:

Refer to Google Cloud's shared responsibility model documentation for more information on OpEx benefits for cloud security.

Question: 372

What is the benefit of using a serverless data processing pipeline service?

A. Pipeline infrastructure is fully managed and scalable.
B. Processed data will not require analysis.
C. Full control over compute resources is provided.
D. Processed data is guaranteed to be free of errors.

Answer: A

Explanation:

Question: 373

An organization has recently completed a migration from on-premises to Google Cloud. How has cost management been affected?

A. Costs will primarily shift from CapEx to OpEx.
B. Costs will primarily shift from OpEx to CapEx.
C. Cost management will stay the same, but the total cost of ownership (TCO) will be lower.
D. Cost management will stay the same, but the total cost of ownership (TCO) will be higher

Answer: A

Explanation:

Question: 374

An organization wants to use an open source library with a flexible ecosystem of tools to create and train its own machine learning models Which product or solution should the organization use1?

A. Cloud Functions
B. TensorFlow
C. Apache Beam
D. Dataflow

Answer: B

Explanation:

Question: 375

An organization has a small development team that has created a web application which runs in a single container. They need a simple serverless and scalable way to host their container Which Google service

should the organization use?

A. Cloud Run
B. Kubernetes Engine
C. App Engine
D. Compute Engine

Answer: A

Explanation:

Question: 376

An organization wants to collect and store all logs generated by applications running in Google Cloud. Which service should they use1?

A. Cloud Logging
B. Cloud Trace
C. Cloud Monitoring
D. Cloud Profiler

Answer: A

Explanation:

Question: 377

An organization needs to increase the speed at which they can train machine learning models Which domain-specific hardware is designed for this task?

A. Preemptible or Spot VMs
B. Containers
C. Bare Metal Solution
D. Cloud TPUs

Answer: D

Explanation:

Question: 378

An organization stores its important industry data in a relational database They want to create a new revenue stream by enabling third parties to use that data in their applications Which cloud first approach should the organization choose?

A. Offer chargeable downloads of archived data.
B. Transfer data into a non-relational database

C. Add third-party users to their database
D. Expose data through a chargeable API

Answer: D

Explanation:

Question: 379

How does the legal concept of data sovereignty affect data*?

A. An individual has the right to control their personal data
B. Data must always be encrypted in transit and at rest.
C. Data is subject to the laws and regulations of the country where it resides.
D. A country has the right to access the data generated within its borders

Answer: C

Explanation:

Question: 380

An organization has hired a team of data scientists and developers They want to create unique value in their business by coding an advanced machine learning model in Vertex Al Workbench Which service should the organization use to train the model?

A. AutoML
B. Prebuilt APIs
C. Custom training
D. Compute Engine

Answer: A

Explanation:

Question: 381

An organization must verify the identity of a user seeking access to a system Which aspect of cloud identity management does this action describe?

A. Auditing
B. Authentication
C. Authorization
D. Encrypting

Answer: B

Explanation:

Question: 382

An organization needs to rapidly scale its use of computing resources and honor its commitment to environmental sustainability. What should the organization do?

A. Purchase energy-efficient servers for an existing on-premises data center.
B. Refactor application software to use less energy
C. Use a carbon-neutral energy provider for an existing on-premises data center
D. Use a public cloud provider with energy-efficient data centers

Answer: D

Explanation:

Question: 383

An organization wants to deploy new workloads to the cloud but must keep some systems onpremises for compliance reasons Both environments must be managed centrally Which type of environment should the organization use1?

A. Hybrid cloud environment
B. Multi cloud environment
C. Container environment
D. Virtual machine environment

Answer: A

Explanation:

Question: 384

Which scenario uses machine learning to unlock business value from unstructured data?

A. Recommending new products based on previous purchases
B. Analyzing tabular records of product defects to predict future maintenance cycles
C. Monitoring financial transactions to identify potential fraud and risk
D. Determining customer sentiments from call center voice recordings

Answer: D

Explanation:

Question: 385

An organization wants to leverage cloud technologies but is concerned about vendor lodc-in.

What would mitigate this concern?

A. Service level agreements
B. Open standards
C. Database services
D. Scalable infrastructure

Answer: A

Explanation:

Question: 386

An organization is using three cloud vendors to maximize their available deployment locations worldwide They are using GKE Enterprise to deploy Kubernetes applications across different clouds What type is this deployment?

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

Answer: C

Explanation:

Question: 387

In Google's cloud security model, how does availability contribute to a robust security posture for data?

A. By checking that data is accurate and trustworthy
B. By ensuring data is reliable and accessible
C. By ensuring data meets industry standards
D. By restricting data access to authorized users

Answer: B

Explanation:

Question: 388

An organization wants to duplicate critical system components to enhance reliability and mitigate single points of failure Which design consideration should the organization use?

A. Latency
B. Security

C. Backups
D. Redundancy

Answer: D

Explanation:

Question: 389

An organization is running SQL Server on-premises and is struggling with capacity and management overhead. They want to modernize this database quickly by using Google products or services What should the organization do?

A. Migrate all SQL Server data to BigQuery
B. Refactor applications to use a cloud first database like Firestore
C. Perform a managed database migration to Cloud SQL
D. Export old tables from SQL Server to Cloud Storage

Answer: C

Explanation:

Question: 390

An organization wants to build custom machine learning models They require a managed platform that provides services to gather data build models and then deploy and monitor those models Which service should they use?

A. Document Al
B. Natural Language API
C. Kubernetes Engine
D. Vertex Al

Answer: D

Explanation:

Question: 391

Which scenario is a good use case for machine learning?

A. Creating customer recommendations
B. Solving ethical dilemmas
C. Classifying data with no prior examples
D. Tasks that require human experience and intuition

Answer: A

Explanation:

Question: 392

An organization has petabytes of data gathered from a wide range of sources They want to use the data for strategic analysis and to guide business decisions What type of service should they use*?

A. Non-relational document database
B. Object store
C. Data warehouse
D. Relational transactional database

Answer: C

Explanation:

Question: 393

An organization wants to access a software application from a cloud vendor without the need to manage their own servers or write their own code Which service model does this represent?

A. Functions as a service
B. Platform as a service
C. Infrastructure as a service
D. Software as a service

Answer: D

Explanation:

Question: 394

An organization wants a centralized view of their cloud infrastructure in a fully managed system that includes uptime checks Which Google Cloud service should they use?

A. Cloud Trace
B. Cloud Monitoring
C. Cloud Logging
D. Cloud Profiler

Answer: B

Explanation:

Question: 395

How does the resource hierarchy in Google Cloud help organizations implement security policies?

A. Policies can be applied at the folder level and are inherited by projects inside the folder
B. Policies can only be applied to individual projects
C. Projects in the resource hierarchy are not affected by security policies.
D. Policies can only be applied at the organization level and affect all projects

Answer: A

Explanation:

Question: 396

An organization needs a flexible and scalable NoSQL database with strong web and mobile application support. Which Google Cloud product or service should the organization use?

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

Answer: C

Explanation:

Question: 397

What is the benefit of using a unified cloud data solution?

A. Data will always cost less to store
B. Data can be automatically secured from external threats
C. Data can enable innovation because it's no longer siloed.
D. Data will always be backed up and cannot be lost or deleted

Answer: C

Explanation:

Question: 398

An organization wants a purpose-built Al solution to increase efficiency and provide personalized interactions for their customer care team Which Google Cloud Al solution should they use?

A. Contact Center Al
B. Cloud Talent Solution
C. Document Al

D. Text-to-SpeechAPI

Answer: C

Explanation:

Question: 399

An organization is concerned that one of their applications takes too long to return a result According to Google's "Four Golden Signals" which signal measures this aspect of the applications performance?

A. Traffic
B. Saturation
C. Errors
D. Latency

Answer: D

Explanation:

Question: 400

What is a defining feature of a non-relational database?

A. Repotting across multiple data sources
B. Queries that join multiple tables
C. A strictly enforced schema
D. A flexible data model

Answer: D

Explanation:

Question: 401

What is the typical cloud spending behavior of most organizations?

A. Decentralized and variable
B. Centralized and variable
C. Decentralized and fixed
D. Centralized and fixed

Answer: A

Explanation:

Question: 402

An organization must identify and fix security vulnerabilities in its cloud infrastructure and applications Which Google Cloud service should they use?

A. VPC networks
B. Security Command Center
C. Google Cloud Armor
D. Cloud Storage

Answer: A

Explanation:

Question: 403

An organization wants to migrate a workload to the cloud without changing the application code or architecture Which migration path describes this approach?

A. Replatformed
B. Rehosted
C. Reimagined
D. Refactored

Answer: B

Explanation:

Question: 404

An organization is running Kubernetes applications across multiple cloud environments They want a consistent and centralized management platform Which service should they choose?

A. Cloud Run
B. Compute Engine
C. GKE Enterprise
D. Cloud Functions

Answer: D

Explanation:

Question: 405

What does the shift toward cloud computing represent for an organization's transformation?

A. An opportunity that is only relevant to the IT department
B. An opportunity to continue business as usual with new cost savings
C. An opportunity to redefine existing business processes and services
D. An opportunity that is limited to large enterprises

Answer: C

Explanation:

Question: 406

An organization wants to transfer some of its data from Google Cloud Which of these statements is true?

A. A technical support ticket must be raised with the correct department
B. Customers have full control of their data and may transfer it at any time
C. Outgoing data transfer must be enabled in the Google Cloud console.
D. Customer data may not be transferred out of Google Cloud

Answer: B

Explanation:

Question: 407

An organization stores backup files in Cloud Storage The files will be accessed annually to test the disaster recovery plan s. Which storage class is the most cost-effective?

A. Nearline class
B. Standard class
C. Coldline class
D. Archive class

Answer: D

Explanation:

Question: 408

An organization is transforming their raw data into a format that can be used to derive business insights Which step of the data value chain does this action represent?

A. Data collection
B. Data analysis
C. Data processing
D. Data storage

Answer: C

Explanation:

Question: 409

An organization is concerned about the unlikely event that Google Cloud infrastructure is physically