Google

Exam Questions Cloud-Digital-Leader

Google Cloud Digital Leader exam

NEW QUESTION 1
- (Topic 1)
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 environment
B. Use labels to segregate resources by environment.
C. Create a single project for all environment
D. Use tags to segregate resources by environment.
E. Create one project for the development environment and one project for the production environment.
F. 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.

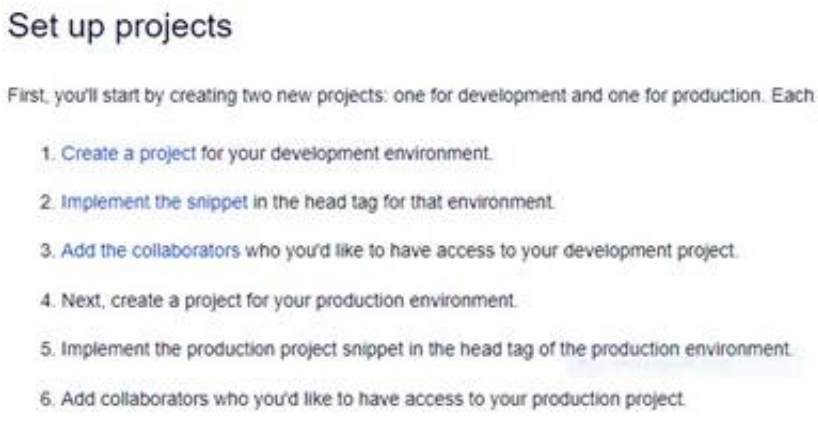

Reference link- https://support.optimizely.com/hc/en-us/articles/4410284353805-Set-up-projects-for-development-and-production-environments

NEW QUESTION 2
- (Topic 1)
Your organization needs to allow a production job to have access to a BigQuery dataset. The production job is running on a Compute Engine instance that is part of an instance group.
What should be included in the IAM Policy on the BigQuery dataset?

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

Answer: C

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

NEW QUESTION 3
- (Topic 1)
Your organization needs to ensure that the Google Cloud resources of each of your departments are segregated from one another. Each department has several environments of its own: development, testing, and production. Which strategy should your organization choose?

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

Answer: B

Explanation:

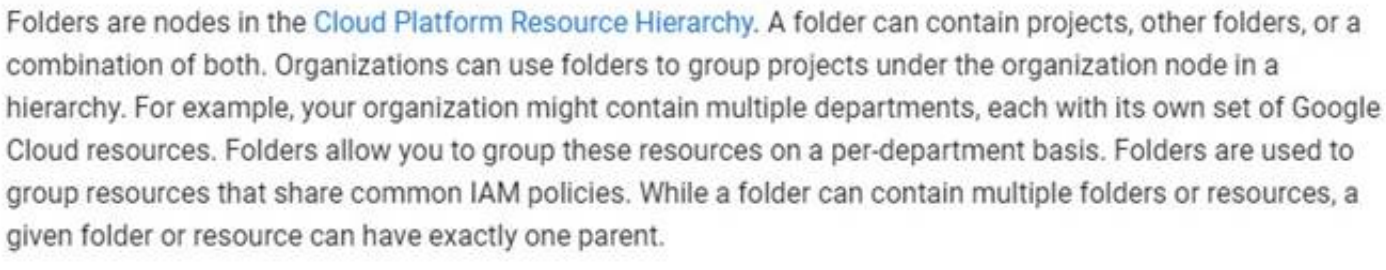

```
# Template for new folder & new project

folder_resource = {
  'name': 'new-folder',
  'type': 'gcp-types/cloudresourcemanager-v2:folders',
  'properties': {
     'parent': 'organizations/99999',
     'displayName': 'new-folder'
  }
}

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

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

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

NEW QUESTION 4
- (Topic 1)
Your organization needs to minimize how much it pays for data traffic from the Google network to the internet. What should your organization do?

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

Answer: A

Explanation:

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

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

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

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

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

NEW QUESTION 5
- (Topic 1)
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:

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

NEW QUESTION 6
- (Topic 1)
How should a multinational organization that is migrating to Google Cloud consider security and privacy regulations to ensure that it is in compliance with global standards?

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

Answer: A

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

NEW QUESTION 7
- (Topic 1)
Your application has repeated data requests of the exact same nature. At the same time, the number of user requests is increasing. Monitoring indicates that the load on the existing database is increasing, and there seems to be a bottleneck. An analysis of the data requested shows us that it is application-managed data and that it changes, but not often. How can you improve the efficiency of the application?

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

Answer: A

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

NEW QUESTION 8
- (Topic 1)
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

NEW QUESTION 9
- (Topic 1)
Your organization is building an application running in Google Cloud. Currently, software builds, tests, and regular deployments are done manually, but you want to reduce work for the team. Your organization wants to use Google Cloud managed solutions to automate your build, testing, and deployment process.
Which Google Cloud product or feature should your organization use?

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

Answer: C

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

NEW QUESTION 10
- (Topic 1)

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

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

Answer: D

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

NEW QUESTION 10
- (Topic 1)
Each of the three cloud service models - infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS) - offers benefits between flexibility and levels of management by the cloud provider and the customer.
Why would SaaS be the right choice of service model?

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

Answer: B

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

NEW QUESTION 11
- (Topic 1)
What is the difference between Standard and Coldline storage?

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

Answer: D

Explanation:

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

NEW QUESTION 12
- (Topic 1)
Your organization is running all its workloads in a private cloud on top of a hypervisor. Your
organization has decided it wants to move to Google Cloud as quickly as possible. Your organization wants minimal changes to the current environment, while using the maximim amount of managed services Google offers.
What should your organization do?

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

Answer: B

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

NEW QUESTION 17
- (Topic 1)
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 log
E. Export the logs to BigQuery to analyze them frequently.

Answer: C

Explanation:

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

NEW QUESTION 19
- (Topic 1)
You are a program manager for a team of developers who are building an event-driven application to allow users to follow one another's activities in the app. Each time a user adds himself as a follower of another user, a write occurs in the real-time database.
The developers will develop a lightweight piece of code that can respond to database writes and generate a notification to let the appropriate users know that they have gained new followers. The code should integrate with other cloud services such as Pub/Sub, Firebase, and Cloud APIs to streamline the orchestration process. The application requires a platform that automatically manages underlying infrastructure and scales to zero when there is no activity.
Which primary compute resource should your developers select, given these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 24
- (Topic 1)
Your organization is developing an application that will capture a large amount of data from millions of different sensor devices spread all around the world. Your organization needs a database that is suitable for worldwide, high-speed data storage of a large amount of unstructured data.
Which Google Cloud product should your organization choose?

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

Answer: D

Explanation:

Reference: https://cloud.google.com/bigtable
Cloud Bigtable is a sparsely populated table that can scale to billions of rows and thousands of columns, enabling you to store terabytes or even petabytes of data. A single value in each row is indexed; this value is known as the row key. Bigtable is ideal for storing very large amounts of single-keyed data with very low latency. It supports high read and write throughput at low latency, and it is an ideal data source for MapReduce operations.
Bigtable is exposed to applications through multiple client libraries, including a supported extension to the Apache HBase library for Java. As a result, it integrates with the existing Apache ecosystem of open-source Big Data software.
Bigtable's powerful back-end servers offer several key advantages over a self-managed HBase installation:
Incredible scalability. Bigtable scales in direct proportion to the number of machines in
your cluster. A self-managed HBase installation has a design bottleneck that limits the performance after a certain threshold is reached. Bigtable does not have this bottleneck, so you can scale your cluster up to handle more reads and writes.
Simple administration. Bigtable handles upgrades and restarts transparently, and it automatically maintains high data durability. To replicate your data, simply add a second cluster to your instance, and replication starts automatically. No more managing replicas or regions; just design your table schemas, and Bigtable will handle the rest for you.
Cluster resizing without downtime. You can increase the size of a Bigtable cluster for a few hours to handle a large load, then reduce the cluster's size again—all without any downtime. After you change a cluster's size, it typically takes just a few minutes under load for Bigtable to balance performance across all of the nodes in your cluster.
Graphical user interface, text, application, email Description automatically generated

NEW QUESTION 26
- (Topic 1)
Your organization is developing a plan for migrating to Google Cloud.
What is a best practice when initially configuring your Google Cloud environment?

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

Answer: B

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

NEW QUESTION 28
- (Topic 1)
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 resource
E. 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.
References:
https://cloud.google.com/compute/docs/instances/reserving-zonal-resources https://cloud.google.com/bigquery/docs/reservations-intro

NEW QUESTION 32
- (Topic 1)
Which of the following is/are true about Bare Metal Solutions?

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

Answer: D

Explanation:
Bare Metal Solution for Oracle
Bring your Oracle workloads to Google Cloud with Bare Metal Solution and jumpstart your cloud journey with minimal risk.
- Continue to run any version, any feature set, any database option, and any customizations (patchsets)
- Enterprise-grade deployment platform
- High availability with Oracle RAC
- Works with any application, any Oracle versions
- All your existing investment in tooling and best practices will work as is

NEW QUESTION 35
- (Topic 1)
Your organization is moving an application to Google Cloud. As part of that effort, it needs to migrate the application's working database from another cloud provider to Cloud SQL. The database runs on the MySQL engine. The migration must cause minimal disruption to users. Data must be secured while in transit.
Which should your organization use?

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

Answer: C

Explanation:

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

NEW QUESTION 38
- (Topic 1)
Your large and frequently changing organization's user information is stored in an on- premises LDAP database. The database includes user passwords and group and organization membership.
How should your organization provision Google accounts and groups to access Google Cloud resources?

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

Answer: C

Explanation:
You can run a single instance of Google Cloud Directory Sync to synchronize user accounts and groups to Google Cloud.
Reference: https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-introduction Text
Description automatically generated https://support.google.com/a/answer/106368?hl=en

NEW QUESTION 42
- (Topic 1)

An organization wants to dynamically adjust its application to serve different user needs. What are the benefits of storing their data in the cloud for this use case?

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

Answer: C

Explanation:

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

NEW QUESTION 43
- (Topic 1)
Your organization recently migrated its compute workloads to Google Cloud. You want these workloads in Google Cloud to privately and securely access your large volume of on- premises data, and you also want to minimize latency.
What should your organization do?

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

Answer: C

Explanation:
Graphical user interface, text, application, Word, email

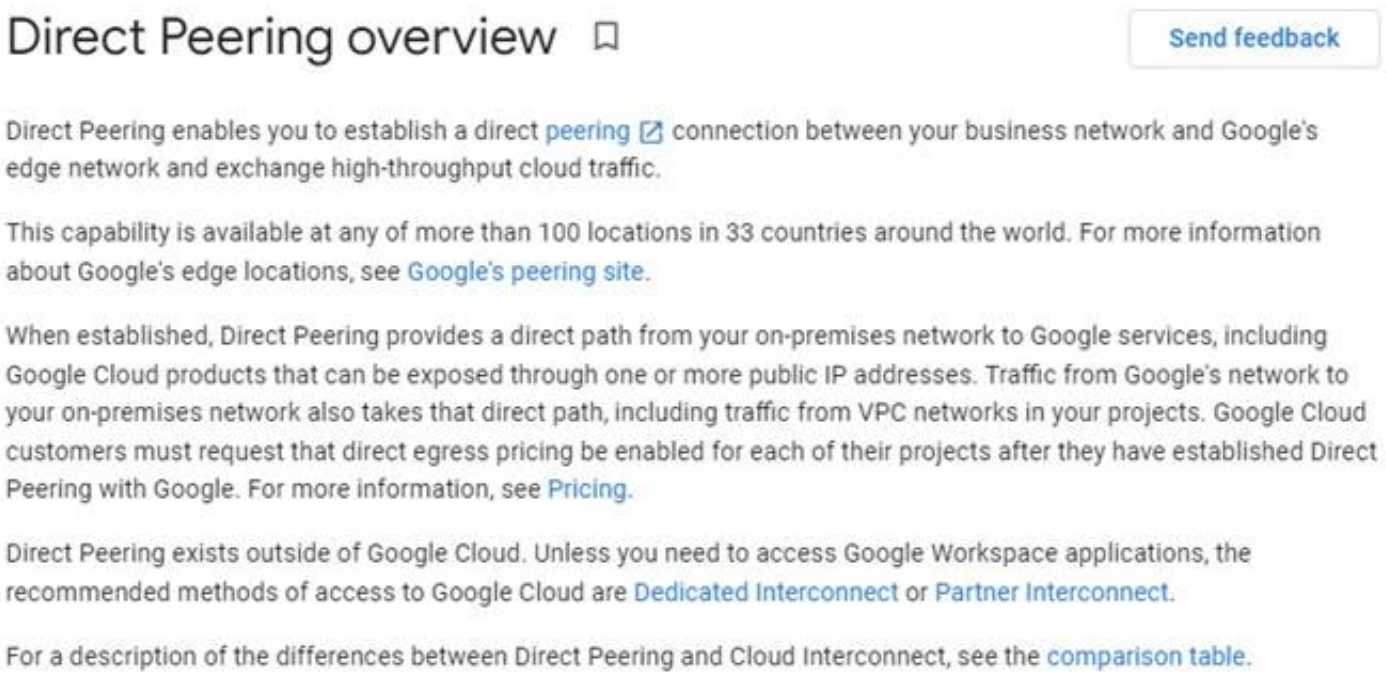

Description automatically generated https://cloud.google.com/network-connectivity/docs/direct-peering

NEW QUESTION 46
- (Topic 1)
Your organization wants to predict the behavior of visitors to its public website. To do that, you have decided to build a machine learning model. Your team has database-related skills but only basic machine learning skills, and would like to use those database skills.
Which Google Cloud product or feature should your organization choose?

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

Answer: A

Explanation:

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

NEW QUESTION 49
- (Topic 1)
What are the key features of Google Cloud Identity.

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

Answer: D

Explanation:
Cloud Identity:
A unified identity, access, app, and endpoint management (IAM/EMM) platform.
- Give users easy access to apps with single sign-on.
- Multi-factor authentication protects user and company data.
- Endpoint management enforces policies for personal and corporate devices
KEY FEATURES :
Modernize IT and strengthen security Multi-factor authentication (MFA)
Help protect your user accounts and company data with a wide variety of MFA verification methods such as push notifications, Google Authenticator, phishing-resistant Titan Security Keys, and using your Android or iOS device as a security key.
Endpoint management
Improve your company's device security posture on Android, iOS, and Windows devices using a unified console. Set up devices in minutes and keep your company data more secure with endpoint management. Enforce security policies, wipe company data, deploy apps, view reports, and export details.
Single sign-on (SSO)
Enable employees to work from virtually anywhere, on any device, with single sign-on to thousands of pre-integrated apps, both in the cloud and on-premises.
Works with your favorite apps
Cloud Identity integrates with hundreds of cloud applications out of the box—and we're constantly adding more to the list so you can count on us to be your single identity platform today and in the future.

NEW QUESTION 54
- (Topic 1)
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.

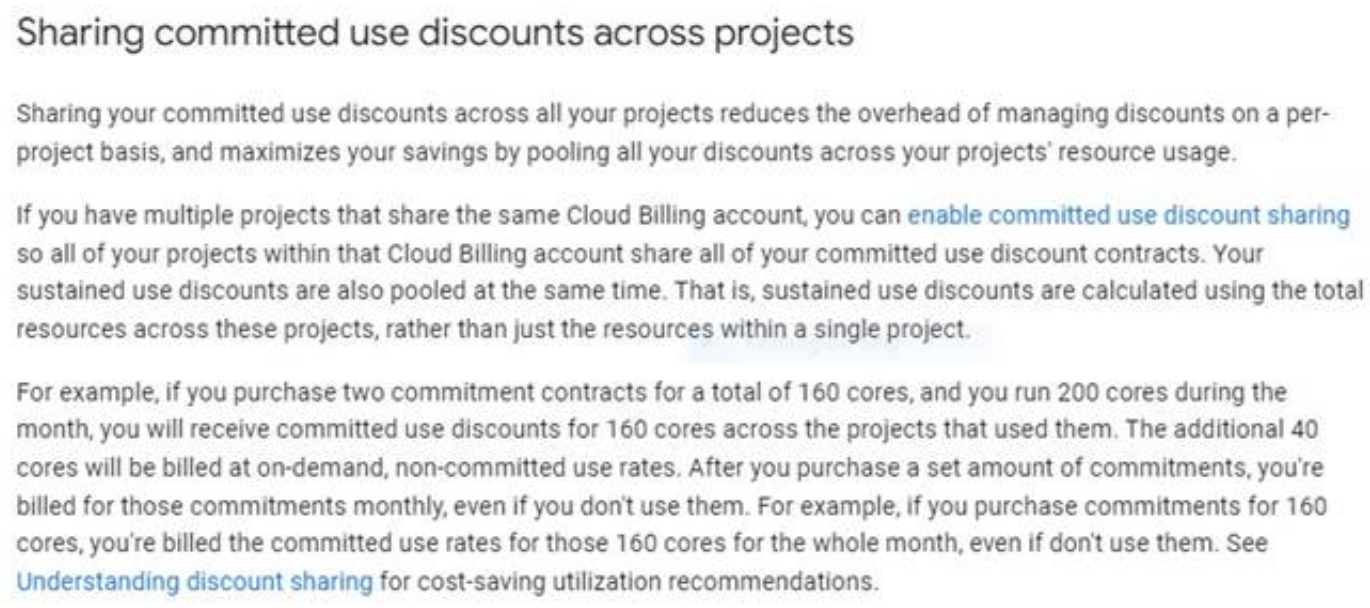
Sharing committed use discounts across projects

Sharing your committed use discounts across all your projects reduces the overhead of managing discounts on a 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.

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 for those 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 for those 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

NEW QUESTION 58
- (Topic 1)
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.

NEW QUESTION 63
- (Topic 1)
What would provide near-unlimited availability of computing resources without requiring your organization to procure and provision new equipment?

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

Answer: A

Explanation:

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

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

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

Answer: B

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

NEW QUESTION 68
- (Topic 1)
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

NEW QUESTION 73
- (Topic 1)
Your company security team manages access control to production systems using an LDAP directory group.
How is this access control managed in the Google Cloud production project?

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

Answer: C

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

NEW QUESTION 76
- (Topic 1)
A partner of yours used to have their own private data center. Your company was already on Google Cloud and now they have also moved to Google Cloud. You are investigating whether there are ways to collaborate better or shared services. What would be one good option to consider?

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

Answer: B

Explanation:
VPC Network Peering allows internal IP address connectivity across two Virtual Private Cloud (VPC) networks regardless of whether they belong to the same project or the same organization.
-> Shared VPC is only within an organization - it allows an organization to connect resources from multiple projects to a common Virtual Private Cloud (VPC) network, so that they can communicate with each other securely and efficiently using internal IPs from that network.
-> Private Google Access is only to access Google APIs and services
References:
-> https://cloud.google.com/vpc/docs/vpc-peering
-> https://cloud.google.com/vpc/docs/private-google-access
-> https://cloud.google.com/vpc/docs/shared-vpc

NEW QUESTION 79
- (Topic 1)
Which Google Cloud service or feature lets you build machine learning models using Standard SQL and data in a data warehouse?

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

Answer: A

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

NEW QUESTION 83
- (Topic 1)
Your team has developed a machine learning model for your customer. The test results indicate very strong predictive capability. The model is then deployed in production. Evaluation of the predictions in production show that they are off by a pronounced margin. What is the issue and how can you solve for it?

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

Answer: D

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

NEW QUESTION 88

- (Topic 3)
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.

NEW QUESTION 92
- (Topic 3)
A large retail organization uses traditional technology for their ecommerce website During peaks m traffic resources are often underutilized or overprovisioned
They have decided to migrate to cloud technology
What aspect of cloud technology will benefit their ecommerce business?

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

Answer: A

NEW QUESTION 96
- (Topic 3)
An organization is training a machine learning model to predict extreme weather events in their country.
How should they collect data to maximize prediction accuracy?

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

Answer: A

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

NEW QUESTION 100
- (Topic 3)
An organization's developers are growing increasingly frustrated by the limitations of their on-premises infrastructure.
How would they benefit from leveraging cloud technology?

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

Answer: C

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

NEW QUESTION 105
- (Topic 3)
Why is data stored in Google Cloud secure and private?

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

Answer: C

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

NEW QUESTION 108
- (Topic 3)
An organization needs to categorize text-based customer reviews on their website using a pre-trained machine learning model.
Which Google Cloud product or service should the organization use?

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

Answer: A

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

NEW QUESTION 110
- (Topic 3)
An organization wants to use BigQuery data analytics to understand their website performance, but wants to move only some data into the cloud.
Which environment should the organization use?

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

Answer: D

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

NEW QUESTION 115
- (Topic 3)
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

NEW QUESTION 118
- (Topic 3)
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

NEW QUESTION 122
- (Topic 3)
An organization is struggling to keep up with the growth of their application which is running on legacy infrastructure.
What might be holding them back?

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

Answer: D

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

NEW QUESTION 123
- (Topic 3)
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

NEW QUESTION 127
- (Topic 3)
An organization's public cloud provider failed to meet their SLA of 99.99% availability.
What is the potential impact on the organization?

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

Answer: C

NEW QUESTION 130
- (Topic 3)
How does a large hotel chain benefit from storing their customer reservation data in the cloud?

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

Answer: C

NEW QUESTION 135
- (Topic 3)
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

NEW QUESTION 139
- (Topic 3)
An organization relies on online seasonal sales for the majority of their annual revenue. Why should the organization use App Engine for their customer app?

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

Answer: B

NEW QUESTION 140
- (Topic 3)
An organization operates their entire IT infrastructure from Google Cloud. What should they do to prepare for data breaches?

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

Answer: C

NEW QUESTION 145
- (Topic 3)
An organization is training a machine learning model to make predictions. What could improve the prediction accuracy of their model?

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

Answer:

C

NEW QUESTION 147
- (Topic 3)
How does switching from on-premises to the cloud help organizations gain value over time?

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

Answer: A

NEW QUESTION 150
- (Topic 3)
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.

NEW QUESTION 154
- (Topic 3)
An organization is struggling to meet user demand for change and wants to modernize their legacy applications by moving the applications to the cloud
Why would this help the organization satisfy user expectations'?

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

Answer: B

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

NEW QUESTION 155
- (Topic 3)
An online retail organization wants to optimize their service.
What is an example of unstructured data that they can use to make decisions?

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

Answer: A

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

NEW QUESTION 159
- (Topic 3)
A retail company stores their product inventory in a legacy system. Often, customers find products on the company's website and want to purchase them in-store. However, when they arrive, they discover that the products are out of stock.
How could the company benefit from using an application programming interface (API)?

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

Answer: D

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

NEW QUESTION 161
- (Topic 3)
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

NEW QUESTION 164
- (Topic 3)
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

NEW QUESTION 166
- (Topic 3)
Why do organizations often struggle to scale their on-premises application infrastructure?

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

Answer: B

NEW QUESTION 167
- (Topic 3)
An international bank is looking for a serverless warehouse solution that lets them perform smart analytics
Which Google Cloud product or service should the bank use?

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

Answer: A

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

NEW QUESTION 170
- (Topic 3)
An organization recently launched a virtual customer support agent, generating vast amounts of text and speech data.
Why should they use a cloud data warehouse to interpret this data?

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

Answer: B

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

NEW QUESTION 175
- (Topic 3)
An organization wants to transform multiple types of structured and unstructured data in the cloud from various sources. The data must be readily accessible for analysis and insights.
Which cloud data storage system should the organization use?

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

Answer: D

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

NEW QUESTION 176
- (Topic 3)
A food delivery service needs access to real-time menu information from all partner restaurants. They also need to share customer order information with the restaurants in real time.
What should the organization use?

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

Answer: B

NEW QUESTION 178
- (Topic 3)
An organization wants to use all available data to offer predictive suggestions on their website that improve over time.
Which method should the organization use?

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

Answer: C

NEW QUESTION 179
- (Topic 3)
An organization is migrating their business applications from on-premises to the cloud. How could this impact their operations and personnel costs?

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

Answer: A

NEW QUESTION 182
- (Topic 3)
An organization wants to use Apigee to manage all their application programming interfaces (APIs).
What will Apigee enable the organization to do?

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

Answer: B

Explanation:
Apigee's API Monitoring enables you to track your APIs to make sure they are up and running correctly. API Monitoring provides near real-time insights into API traffic and performance, to help you quickly diagnose and solve issues as they arise.
Apigee works with APIs not necessarily applications. It allows organizations to gain actionable insights across the entire API value chain and monetize API products and maximize the business value of digital assets. https://cloud.google.com/apigee#section-11

NEW QUESTION 184
- (Topic 3)
An organization wants to move from a tactical cloud adoption approach to a transformational approach.
How should they adapt the way they lead the organization?

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

Answer: A

NEW QUESTION 189
- (Topic 3)
An organization wants to introduce a new image recognition login system What should the organization do to follow SRE principles?

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

Answer: A

NEW QUESTION 190
- (Topic 3)
How does Cloud SQL help organizations create business insights?

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

Answer: A

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

NEW QUESTION 192
- (Topic 3)
An organization meets their service level objective (SLO) of 99 999% ("five nines") How much downtime do their end users experience per year?

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

Answer: A

NEW QUESTION 194
- (Topic 2)
The customer has applications that do data processing on-premise. They have been built using Ha-doop and Spark. What product should I use on Google Cloud?

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

Answer: A

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

NEW QUESTION 198
- (Topic 2)
If you increase the size of a subnet in a custom VPC network, the IP addresses of virtual machines already on that subnet might be affected. Which options are Correct.

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

Answer: A

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

NEW QUESTION 199
- (Topic 2)
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.

NEW QUESTION 201
- (Topic 2)
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

NEW QUESTION 204
- (Topic 2)
A financial services company is running an experimental application workload that has a very large number of mathematical calculations involving floating-point numbers. The current application that is running on compute engine is not providing enough speed and throughput. What are the options to increase the processing performance?

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

Answer: D

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

NEW QUESTION 208
- (Topic 2)
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.

NEW QUESTION 212
- (Topic 2)
Which of the followings are core components of Anthos?

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

Answer: D

Explanation:

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

NEW QUESTION 214
- (Topic 2)
What load balancer type is supported with Cloud Armor security policies?

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

Answer: D

Explanation:
Google Cloud Armor security policies protect your application by providing Layer 7 filtering and by scrubbing incoming requests for common web attacks or other Layer 7 attributes to potentially block traffic before it reaches your load balanced backend services or backend buckets. Each security policy is made up of a set of rules that filter traffic based on conditions such as an incoming request's IP address, IP range, region code, or request headers.
-> Google Cloud Armor security policies are available only for backend services behind an external HTTP(S) load balancer. The load balancer can be in Premium Tier or Standard Tier.
-> Google Cloud Armor security policies and IP DENY lists and ALLOW lists are available only for HTTP(S) load balancing.
Reference link- https://cloud.google.com/armor/docs/security-policy-overview

NEW QUESTION 219
- (Topic 2)
A customer has contacted you about migrating to Google Cloud. The customer would like to mi-grate their data from on premises as soon as possible. They don't have the budget to rewrite code, and they want the most direct route. What migration option should suggest to the customer?

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

Answer: C

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

Lift and shift

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

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

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

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

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

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

NEW QUESTION 221
- (Topic 2)
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.

NEW QUESTION 223
- (Topic 2)
Which of the following statements describe the features of a preemptible VM in-stance? (Select Three Answer)

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

Answer: CDE

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

NEW QUESTION 228
- (Topic 2)
Which of the following is/are core storage options available on the Google Cloud Platform?

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

Answer: D

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

services to get the job done.

NEW QUESTION 233
- (Topic 2)
While on-premise, an enterprise had multiple teams, each with its own analytics data store. Attempts to converge the storage for centralized, company-wide analysis failed because of speed and scaling issues. What would be the preferred destination architecture on Google Cloud?

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

Answer: C

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

NEW QUESTION 238
- (Topic 2)
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.

NEW QUESTION 243
- (Topic 2)
Your Coustomer's Organization has decided to move to the cloud. They currently run VMs on-promise but their goal on Google cloud is to run containers, primarily on Google Kuber- nete's Engine. They have a lease for their private data center for another year that they have already paid for. What could be strategy they could adopt in migrating?

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

Answer: B

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

NEW QUESTION 247
- (Topic 2)
You are a program manager in a company and handling a project and you need to create a virtual machine on google cloud console that will be very simple to set up, by flipping a bit via command, API, or with developer console that gives you 30 seconds to shut down when you’re preempted, allow you to save your work that also helps in the company budget upto 70-80% of less charges than the regular VMs.

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

Answer: B

Explanation:
Preemptible VMs have all these features
Simple configuration
Create a preemptible instance simply by flipping a bit via command, API, or developer console.
Easy extensibility
Attach GPUs and local SSDs to preemptible instances for additional performance and savings.

Graceful shutdown
Compute Engine gives you 30 seconds to shut down when you're preempted, letting you save your work in progress for later.
Large scale computing
Spin up as many instances as you need and turn them off when you're done. You only pay for what you use.
Quickly reclaim capacity
Managed instance groups automatically recreate your instances when they're preempted (if capacity is available).
Fixed pricing
Preemptible VMs have fixed pricing up to 80% off regular instances. They show up on your bill separately so you'll see just how much you're saving.

NEW QUESTION 249
- (Topic 2)
You have deployed a new public web application that allows users to register and login with email ids, phone numbers, or user ids. You are seeing some unusual activity with user registrations and logins from a few IPs. A large number of accounts were created very quickly. Logins are also hap-pening quickly thereafter from these new accounts. Different parts of the application are being ex-plored, all of which are putting a heavy load on the application. What could be a problem and how can you solve it?

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

Answer: B

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

NEW QUESTION 253
- (Topic 2)
Which of the following is true while creating a boot persistent disk from a snapshot.

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

Answer: D

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

NEW QUESTION 254
- (Topic 2)
You are consulting for a client who is migrating to Google Cloud. They presently have a matrix or-ganization. Their IT environments were managed around projects. Each team had multiple projects. All the projects had a flat structure under the company. What would you advise them when plan-ning for the move?

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

Answer: A

Explanation:
Folders for a related group of projects are the recommended approach.
-> A flat structure under the organization node is possible on Google Cloud, but it is not recommended. It becomes tougher to manage.
-> Projects cannot have sub-projects; there can only be resources within Projects.
Reference link- https://cloud.google.com/resource-manager/docs/cloud-platform- resource-hierarchy

NEW QUESTION 256
- (Topic 2)
In discussions with a prospective customer who wants to move to Google Cloud to make use of the latest, scalable technologies available therein, you learn that there are very strict regulations concern-ing the storage of data. They only have the approval to store it in their current private data cen-ter. What would you advise them?

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

Answer: A

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

NEW QUESTION 261
- (Topic 2)
In Google Cloud IAM: if a policy applied at the project level gives you Owner permissions, your access to an individual resource in that project might be restricted to View permission if someone applies a more restrictive policy directly to that resource. What is correct below the options

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

Answer: A

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

NEW QUESTION 263
- (Topic 2)
You are running a data warehouse on BigQuery. A partner company is offering a recommendation engine based on the data in your data warehouse. The partner company is also running their applica-tion 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 da-taset. 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:
- 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

NEW QUESTION 267
- (Topic 2)
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 conver-sation 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

NEW QUESTION 272
- (Topic 2)
What characteristics should an organization adopt to be a DevOps organization?

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

Answer: D

Explanation:

NEW QUESTION 274
- (Topic 2)
What according to you are NOT the key capabilities of In-App Messaging?

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

Answer: C

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

NEW QUESTION 276
- (Topic 2)
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.

NEW QUESTION 281
- (Topic 2)
Cloud Data Loss Prevention (DLP) is a fully managed service designed to help discover, classify, and protect the most sensitive data. 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

NEW QUESTION 284
- (Topic 2)
A customer has an application running in virtual machines. They are migrating this application to Google Cloud. They have previously had scaling issues when on-premises as VMs had to be pre-allocated. Capacity planning was repeatedly off mark - it's either too many VMs or too less. They want to match the capacity to demand while keeping the application running always. They don't have the time or budget to re-architect the systems using containers and Kubernetes at the moment. What would be your recommendation?

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

Answer: B

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

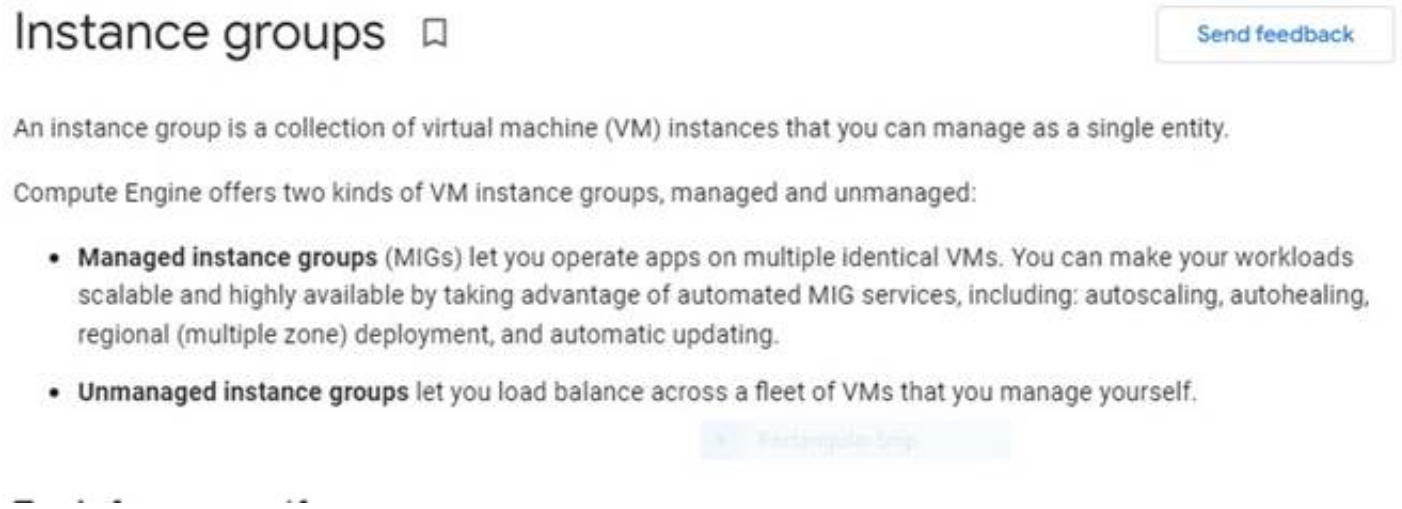

https://cloud.google.com/compute/docs/instance-groups

NEW QUESTION 286
- (Topic 2)
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

NEW QUESTION 288
- (Topic 2)
Customer Managed Encryption Keys (CMEK) can be used for encrypting data inside Cloud BigTable, which of the following statements is/are correct. (Select two answer)

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

Answer: BC

Explanation:
Customer-managed encryption keys for Cloud BigTable.
By default, all the data at rest in Cloud Bigtable is encrypted using Google's default encryption. Bigtable handles and manages this encryption for you without any additional action on your part.

If you have specific compliance or regulatory requirements related to the keys that protect your data, you can use customer-managed encryption keys (CMEK) for BigTable. Instead of Google managing the encryption keys that protect your data, your BigTable instance is protected using a key that you control and manage in Cloud Key Management Service (Cloud KMS).
Features
Security: CMEK provides the same level of security as Google's default encryption but provides more administrative control.
Data access control: Administrators can rotate, manage access to, and disable or destroy the key used to protect data at rest in BigTable .
Auditability: All actions on your CMEK keys are logged and viewable in Cloud Logging. Comparable performance: BigTable CMEK-protected instances offer comparable performance to BigTable instances that use Google default encryption.
Flexibility: You can use the same CMEK key in multiple projects or instances or you can use separate keys, depending on your business needs.

NEW QUESTION 289
- (Topic 2)
You are looking for a one stop reference page for GCP support. What Page would you se- lect?

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

Answer: C

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

NEW QUESTION 292
- (Topic 2)
You are a database manager working for a new product that will need millions of reading and writ-ing from the database, with zero downtime, key-value i.e. NoSQL features, no manual steps should be required to ensure consistency, repair data, synchronize writes and deletes, Which of the follow-ing database you choose?

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

Answer: B

Explanation:
Cloud BigTable
Key features
High throughput at low latency
Bigtable is ideal for storing very large amounts of data in a key-value store and supports high read and write throughput at low latency for fast access to large amounts of data. Throughput scales linearly—you can increase QPS (queries per second) by adding Bigtable nodes. Bigtable is built with proven infrastructure that powers Google products used by billions such as Search and Maps.
Cluster resizing without downtime
Scale seamlessly from thousands to millions of reads/writes per second. Bigtable throughput can be dynamically adjusted by adding or removing cluster nodes without restarting, meaning you can increase the size of a Bigtable cluster for a few hours to handle a large load, then reduce the cluster's size again—all without any downtime. Flexible, automated replication to optimize any workload
Write data once and automatically replicate where needed with eventual consistency—giving you control for high availability and isolation of reading and write workloads. No manual steps are needed to ensure consistency, repair data, or synchronize writes and deletes. Benefit from a high availability SLA of 99.999% for instances with multi- cluster routing across 3 or more regions (99.9% for single-cluster instances).

NEW QUESTION 295
- (Topic 2)
What service is a fully managed real-time messaging service that allows you to send and receive messages between independent applications.

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

Answer: B

Explanation:
Google Cloud Pub/Sub is a scalable, durable event ingestion and delivery system.
-> Pub/Sub allows services to communicate asynchronously, with latencies on the order of 100 milliseconds.
-> Pub/Sub is used for streaming analytics and data integration pipelines to ingest and distribute data. It is equally effective as messaging-oriented middleware for service integration or as a queue to parallelize tasks.
-> Pub/Sub enables you to create systems of event producers and consumers, called publishers and subscribers. Publishers communicate with subscribers asynchronously by broadcasting events, rather than by synchronous remote procedure calls (RPCs).
Reference link- https://cloud.google.com/pubsub/docs/overview

NEW QUESTION 297
- (Topic 2)
You are working in a company where you need to store Terabytes of Image Data daily and process them e.g. Taking photos of the entire planet 24 hours every day with satellite and sending data to data centres to store and process it. Which of the following would be the best combination for your infrastructure.
You are working in a company where you need to store Terabytes of Image Data daily and process them e.g. Taking photos of the entire planet 24 hours every day with satellite and

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

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

Answer: C

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

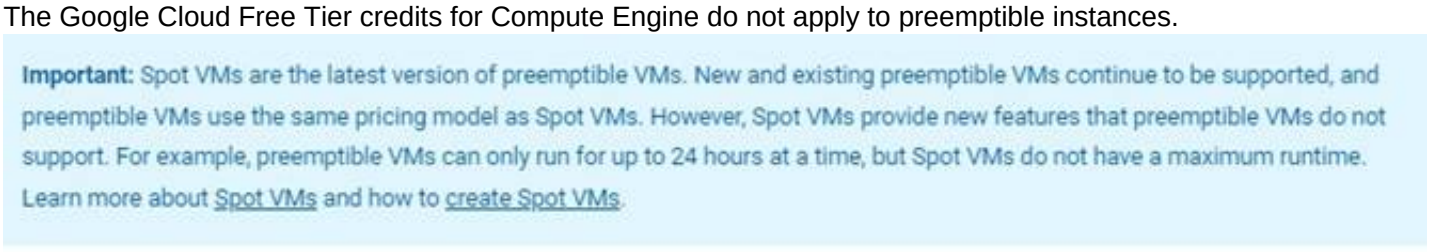

Reference link- https://cloud.google.com/compute/docs/instances/preemptible

NEW QUESTION 300
- (Topic 2)
You have experimented with Google Cloud using your own credit card and expensed the costs to your company. Your company wants to streamline the billing process and charge the costs of your projects to their monthly invoice. What should you do?

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

Answer: B

Explanation:
To change the Cloud Billing account for a project, you need to be able to move a project from one Cloud Billing account to another. To accomplish this task, you need permissions adequate to unlink the project from the existing Cloud Billing account AND to link the project to the target Cloud Billing account. Roles with adequate permissions to perform this task: Project Owner or Project Billing Manager on the project, AND Billing Account Administrator or Billing Account User for the target Cloud Billing account
interface, text, application, email Description automatically generated
Reference link- https://cloud.google.com/billing/docs/how-to/modify-

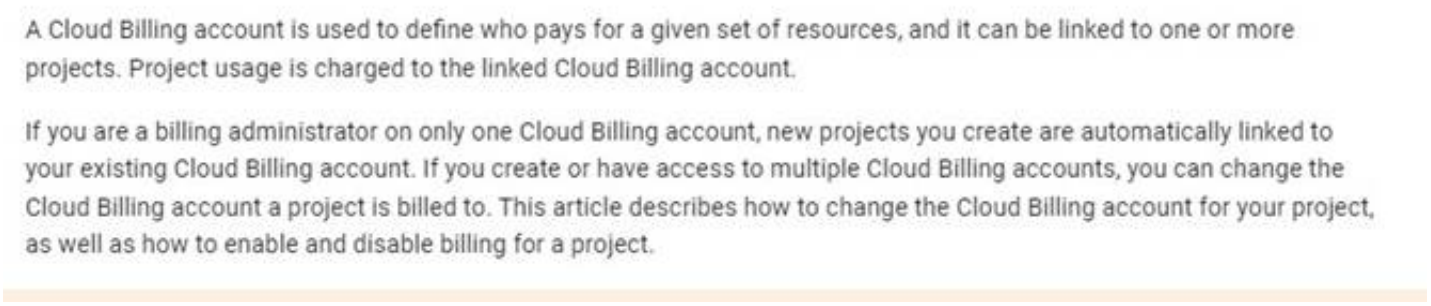

NEW QUESTION 303
- (Topic 2)
An organization wants to evaluate the performance of their entire cloud infrastructure, including metrics like server uptime and response rate reports. Which Google Cloud tool should the organi-zation use?

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.

NEW QUESTION 306
- (Topic 2)
You have a well established development and operations team. Your teams were managing the en-tire software delivery/deployment cycle on-premise. When migrating to the cloud, you want to con-tinue having this approach. Which is the ideal option for you?

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

Answer: D

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

NEW QUESTION 307
- (Topic 2)
You are a DevOps Engineer in an E-commerce company that sells products globally, across the countries, Customers buy products, add them to carts or check-in stock from different parts of the world with different timestamps, you need to choose a database that can scale globally without any hassle and lots of developer support, it should be consistent across regions, can scale horizontally to support enormous user, automatically replicates, shards and even auto transaction processing. Which of the following database do you choose?

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

Answer: B

Explanation:
Cloud Spanner:
Fully managed relational database with unlimited scale, strong consistency, and up to 99.999% availability.
- Get all the benefits of relational semantics and SQL with unlimited scale
- Start at any size and scale with no limits as your needs grow
- Enjoy high availability with zero scheduled downtime and online schema changes
- Deliver high-performance transactions with strong consistency across regions and continents
- Focus on innovation, eliminating manual tasks with capabilities like automatic sharding
Automatic sharding
Cloud Spanner optimizes performance by automatically sharding the data based on request load and size of the data. As a result, you can spend less time worrying about how to scale your database and instead focus on scaling your business.
Strong transactional consistency
Purpose-built for external, strong, global transactional consistency.
Regional and multi-regional configurations
No matter where your users may be, apps backed by Cloud Spanner can read and write up-to-date strongly consistent data globally Additionally, when running a multi-region instance, your database is able to survive a regional failure, and offers industry-leading 99.999% availability.
Online schema changes with no downtime
Cloud Spanner users can make a schema change, whether it's adding a column or adding an index while serving traffic with zero downtime. Hence you now have the flexibility to adapt your database to your business needs without compromising on the availability of your application.

NEW QUESTION 309
- (Topic 2)
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:
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.

NEW QUESTION 312
- (Topic 2)
You are working in a company that provides different services to its customer. Now it also wants to offer some paid API services to its B2B customers for e.g. google provides google maps API, cloud vision API, and language translation API. You need to figure out the best solution for the service.

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

Answer: C

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

NEW QUESTION 314
- (Topic 1)
Your Google Cloud Platform [GCP] admin has to manage a bunch of API keys for external services that are accessed by different applications, which are used by a few teams. What is the best way to manage them?

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

Answer: B

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

NEW QUESTION 315
- (Topic 1)
A fitness band company is continuously ingesting data from millions of its consumers. Different kinds of data based on time, like location, heartbeat rate, temperature, movement, etc. are connect-ed. They need a high throughput database that can write data very fast. Since their users are spread across the world, they need the database to be geographically scalable. Consumers also want to see near-real-time visualizations of their activities. Which of these databases would be a good fit?

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

Answer: B

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

NEW QUESTION 318
- (Topic 1)
Your organization needs to categorize objects in a large group of static images using machine learning. Which Google Cloud product or service should your organization use?

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

Answer: C

Explanation:

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

NEW QUESTION 323
- (Topic 1)
Your company’s development team is building an application that will be deployed on Cloud Run. You are designing a CI/CD pipeline so that any new version of the application can be deployed in the fewest number of steps possible using the CI/CD pipeline you are designing. You need to select a storage location for the images of the application after the CI part of your pipeline has built them.
What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 325
- (Topic 1)
An organization has created an ecommerce website. What data on this website would be considered structured data?

A. Product photographs

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

Answer: D

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

NEW QUESTION 330
- (Topic 1)
Your organization consists of many teams. Each team has many Google Cloud projects. Your organization wants to simplify the management of identity and access policies for these projects.
How can you group these projects to meet this goal?

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

Answer: C

Explanation:

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

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

NEW QUESTION 332
- (Topic 1)
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 logo
E. Write only the code to integrate in-to your workflow.

Answer: B

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

NEW QUESTION 334
- (Topic 1)
What are the network requirements for Private Google Access?

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

Answer: B

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

NEW QUESTION 338
- (Topic 1)
An organization has had a data leak scare because one employee made a sensitive Cloud Storage bucket available to the public. Given the nature of the company's business, it is understood that there is never any reason to give the public direct access to any file. The
security head wants to ensure that such an event never occurs again. How can you ensure this?

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

Answer: B

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

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

NEW QUESTION 342
- (Topic 1)
Your organization needs to analyze data in order to gather insights into its daily operations. You only want to pay for the data you store and the queries you perform. Which Google Cloud product should your organization choose for its data analytics warehouse?

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

Answer: D

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

NEW QUESTION 345
- (Topic 1)
You are currently managing workloads running on Windows Server for which your company owns the licenses. Your workloads are only needed during working hours, which allows you to shut down the instances during the weekend. Your Windows Server licenses are up for renewal in a month, and you want to optimize your license cost.
What should you do?

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

Answer: D

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

NEW QUESTION 349

- (Topic 1)
Which Google Cloud product gives you a consistent platform for multi-cloud application deployments and extends other Google Cloud services to your environment?

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

Answer: D

Explanation:

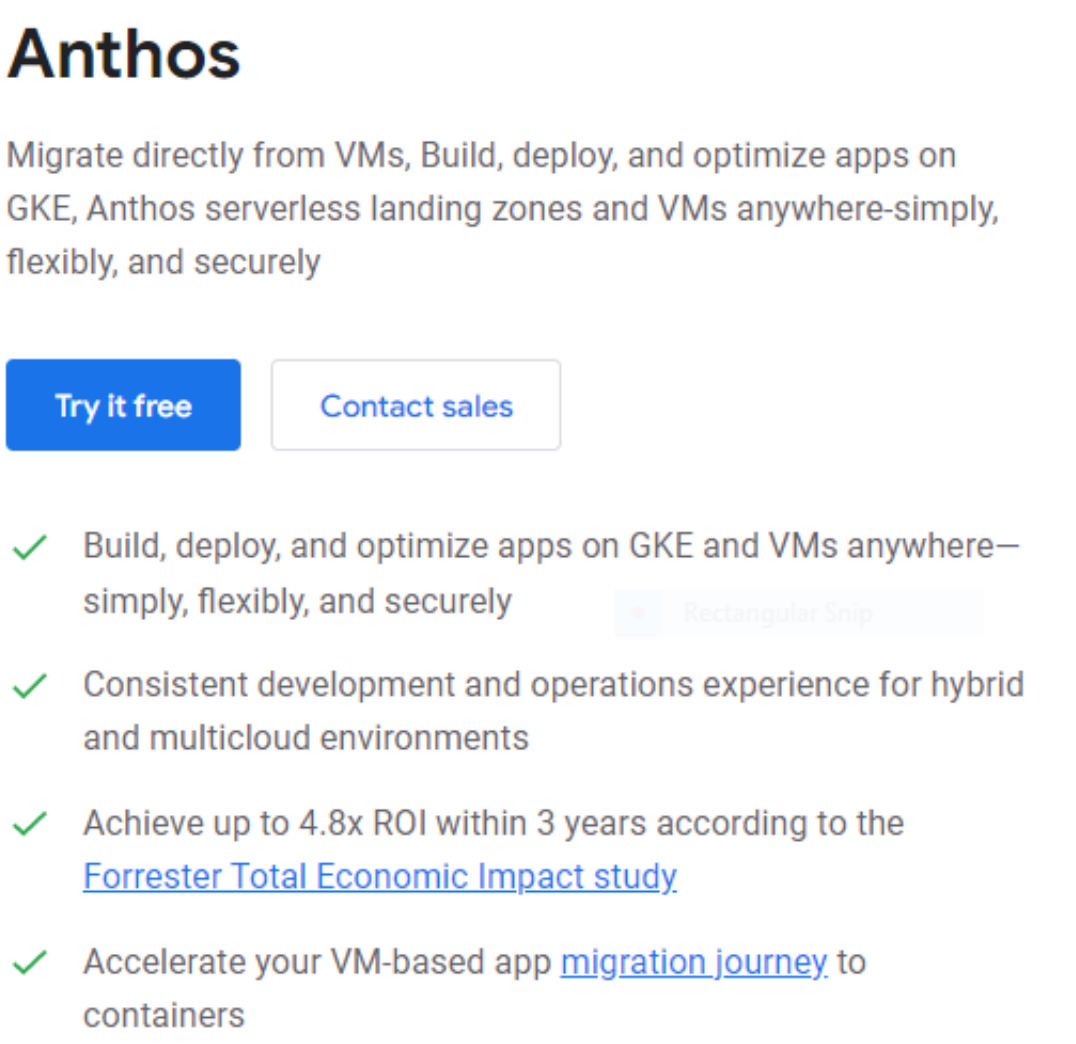

https://cloud.google.com/anthos

NEW QUESTION 353
- (Topic 1)
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.

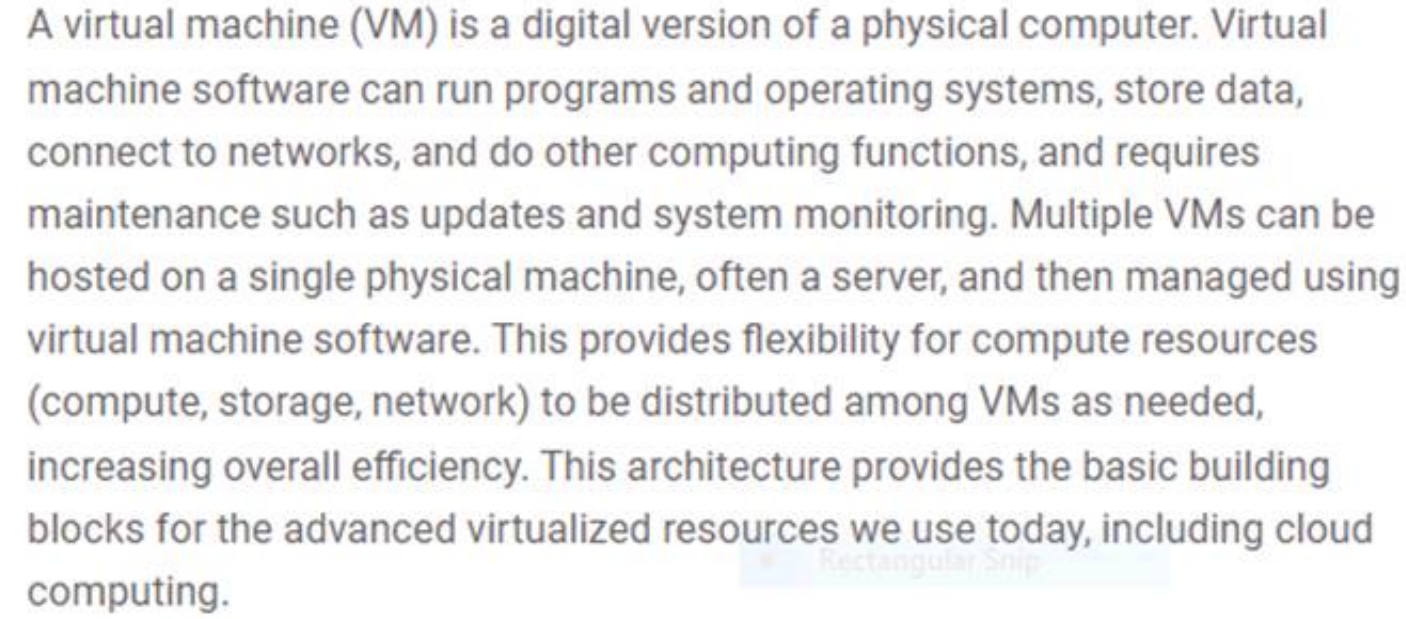

Table
Description automatically generated with medium confidence https://cloud.google.com/learn/what-is-a-virtual-machine

NEW QUESTION 354
- (Topic 1)
A customer has new applications to build that has to handle both batch data and streaming data. Which product should they choose?

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

Answer: B

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

NEW QUESTION 355
- (Topic 1)
There are internal compliance requirements that demand that we do not use any APIs or services that are not backed by SLAs. Which of these are acceptable for us? (Choose two answer)

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

Answer: CD

Explanation:

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

NEW QUESTION 360
- (Topic 1)
Which of the following statements is/are correct about Bare Metal Solutions?

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

Answer: D

Explanation:

Option A is true
You are responsible for the licensing of all of your software. Bare Metal Solution uses a bring-your-own-license (BYOL) model.
Apart from this you are responsible for the software, applications, and data that you use and store in the Bare Metal Solution environment.
Responsibilities Data, including:
- Security and encryption

- Backups
Software and applications, including:
- Installation
- Configuration
- Upgrades and patching
Operating system and any hypervisor, including:
- Configuration changes
- Upgrades and patching Server clusters, including:
- Installation
- Configuration
- Maintenance Licensing
Option B & C is also true.
With Bare Metal Solution, Google Cloud provides and manages the core infrastructure, the net-work, the physical and network security, and hardware monitoring capabilities in an
environment from which you can access all of the Google Cloud services. The core infrastructure includes secure, controlled-environment facilities, and power.
The Bare Metal Solution also includes the provisioning and maintenance of the custom, sole-tenancy hardware with local SAN, and smart hands support.
The network, which is managed by Google Cloud includes a low-latency Cloud Interconnect con-nection into the customer Bare Metal Solution environment.
The available Google Cloud services include private API access, management tools, support, and billing.

NEW QUESTION 364
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

Cloud-Digital-Leader Practice Exam Features:

* Cloud-Digital-Leader Questions and Answers Updated Frequently

* Cloud-Digital-Leader Practice Questions Verified by Expert Senior Certified Staff

* Cloud-Digital-Leader Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* Cloud-Digital-Leader Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year