

## Exam Questions Cloud-Digital-Leader

Google Cloud Digital Leader exam

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



**NEW QUESTION 1**

- (Topic 1)

An organization runs their application on a virtual machine, but every time they want to edit specific features, they have to bring the system offline to update the application. What would be a more appropriate solution for their app?

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

**Answer:** B

**Explanation:**

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

## What are containers?

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

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

**NEW QUESTION 2**

- (Topic 1)

Your organization 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.

### Set up projects

First, you'll start by creating two new projects: one for development and one for production. Each project will need its own snippet:

1. Create a project for your development environment.
2. Implement the snippet in the head tag for that environment.
3. Add the collaborators who you'd like to have access to your development project.
4. Next, create a project for your production environment.
5. Implement the production project snippet in the head tag of the production environment.
6. Add collaborators who you'd like to have access to your production project.

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

**NEW QUESTION 3**

- (Topic 1)

Your company has multiple internal applications used by your employees. You also have to give access to certain vendors and contractors. What is a good option for you to adopt?

- A. Keep the credentials separate for each application to reduce the blast radius in case of any issues.
- B. Use an external identity provider that is famous and popular like Facebook or Twitter; that way, even your vendors and contractors will have an account there.
- C. Allow all users, especially contractors and vendors, to bring their own identities, like those at gmail.com.

D. Use an IDaaS (Identity as a Service) product that can centrally manage authentication and authorization for the applications.

**Answer:** D

**Explanation:**

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

#### NEW QUESTION 4

- (Topic 1)

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

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

**Answer:** D

**Explanation:**

Cloud Identity Platform

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

Reference link- <https://cloud.google.com/identity-platform>

#### NEW QUESTION 5

- (Topic 1)

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

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

**Answer:** C

**Explanation:**

Preemptible VMs:

Predictable and low cost

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

Expand your batch processing

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

rendering/transcoding, and testing.

Get more from your containers

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

Enable it instantly

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

#### 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)

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 maximum 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 10**

- (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 14

- (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 team 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 19

- (Topic 1)

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

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

**Answer: A**

**Explanation:**

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

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

Fully managed Cloud Run is an ideal **serverless platform** for stateless containerized microservices that don't require Kubernetes features like namespaces, co-location of containers in pods (sidecars) or node allocation and management.

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

#### NEW QUESTION 23

- (Topic 1)

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

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

**Answer: C**

**Explanation:**

Graphical user interface, text, application, Word, email

## Direct Peering overview

[Send feedback](#)

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

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

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

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

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

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

**NEW QUESTION 25**

- (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 29**

- (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](https://cloud.google.com/compute/docs/instances/signing-up-committed-use-discounts#sharing_committed_use_discounts_across_projects)

### NEW QUESTION 34

- (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 35

- (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.

# Database Migration Service

Simplify migrations to Cloud SQL. Available now for MySQL and PostgreSQL, with SQL Server migrations and Oracle to PostgreSQL migrations in preview.

[Get started](#)[Migration guide](#)

- ✓ Migrate databases to Cloud SQL from on premises, Google Compute Engine, and other clouds
- ✓ Replicate data continuously for minimal downtime migrations
- ✓ Serverless and easy to set up

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

## NEW QUESTION 39

- (Topic 1)

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

- A. An input variable is used in making prediction
- B. E.
- C. number of rooms in a house price prediction model.
- D. One row of a dataset containing one or more input columns and possibly a prediction result.
- E. An answer for a prediction task, either the answer produced by a machine learning system or the right answer supplied in training data
- F. E.
- G. image contains a "cat".
- H. The "knobs" that you tweak during successive runs of training a model
- I. E.
- J. learning rate

**Answer:** B

### Explanation:

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

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

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

## NEW QUESTION 44

- (Topic 1)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 45**

- (Topic 1)

An organization wants to move from a strategic cloud adoption maturity level to a trans-formational one. How should the organization change the way they scale?

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 47**

- (Topic 1)

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 52**

- (Topic 1)

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 57**

- (Topic 3)

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

- A. Filestore
- B. Cloud Spanner
- C. Data Catalog
- D. Cloud Storage

**Answer:** D

**Explanation:**

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

**NEW QUESTION 60**

- (Topic 3)

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

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

**Answer:** D

#### NEW QUESTION 65

- (Topic 3)

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

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

**Answer:** B

#### NEW QUESTION 68

- (Topic 3)

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

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

**Answer:** A

#### NEW QUESTION 72

- (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 77

- (Topic 3)

An organization needs to migrate specialized workloads to the cloud while maintaining their existing complex licensing and architecture.

What Google Cloud solution should the organization use?

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 82

- (Topic 3)

An organization delivers a proactive healthcare service. They want to efficiently and automatically collect patient data.

What should the organization encourage the patients to do?

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

**Answer:** B

#### NEW QUESTION 85

- (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 90

- (Topic 3)

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

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

**Answer:** D

#### Explanation:

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

#### NEW QUESTION 93

- (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 96

- (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 99

- (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 102

- (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 103

- (Topic 3)

A cloud-native organization is not meeting their service level objective (SLO) but has not exhausted their error budget.

What should the organization prioritize?

- A. Innovation to improve user experience
- B. Hardware reliability to improve availability
- C. Stability to avoid prolonged user downtime
- D. Speed to release new features

**Answer:** C

#### Explanation:

Both Devs and SRE team must ensure that the error budget does not become exhausted. To avoid it, releases have to stop for the time being until the error budget resets. The team would have to reprioritize to focus on reliability to get it back to an acceptable state.

#### NEW QUESTION 104

- (Topic 3)

Which policy helps Google Cloud keep customer data private?

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

**Answer:** B

#### NEW QUESTION 108

- (Topic 3)

An organization wants a cost-effective relational database.

Which Google Cloud service should the organization use?

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

**Answer:** C

#### NEW QUESTION 111

- (Topic 3)

How would a global organization benefit from managing their data with Cloud Spanner?

- A. Cloud Spanner is optimized for cold storage
- B. Cloud Spanner replicates data across regions in real time
- C. Cloud Spanner is optimized to ingest unstructured data
- D. Cloud Spanner visualizes and analyzes data in real time

**Answer:** B

#### Explanation:

Spanner is Google's scalable, multi-version, globally-distributed, and synchronously-replicated database.

#### NEW QUESTION 114

- (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 116

- (Topic 3)

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

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

**Answer: D**

#### NEW QUESTION 120

- (Topic 3)

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

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

**Answer: A**

#### NEW QUESTION 124

- (Topic 3)

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

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

**Answer: B**

#### NEW QUESTION 125

- (Topic 3)

An organization wants its users to validate a series of new features for their app. Why should they use App Engine?

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

**Answer: C**

#### NEW QUESTION 129

- (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 130

- (Topic 3)

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

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

**Answer: C**

#### NEW QUESTION 133

- (Topic 3)

What is artificial intelligence?

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

D. Any system capable of a task that normally requires human cognition

**Answer:** D

**NEW QUESTION 136**

- (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 140**

- (Topic 3)

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

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

**Answer:** B

**Explanation:**

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

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

**NEW QUESTION 145**

- (Topic 3)

What does Cloud Debugger help an organization do?

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 149**

- (Topic 2)

The customer has applications that do data processing on-premise. They have been built using Hadoop 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 150**

- (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 154**

- (Topic 2)

You are working with a user to set up an application in a new VPC behind a firewall and it is noticed that the user is concerned about data egress. Therefore, to provide assistance you want to configure the fewest open egress ports. Which of the following statements is correct?

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

**Answer: B**

**Explanation:**

Implied rules Every VPC network has two implied firewall rules. These rules exist, but are not shown in the Cloud Console:

Implied allow egress rule. An egress rule whose action is allow, destination is 0.0.0.0/0, and priority is the lowest possible (65535) lets any instance send traffic to any destination, except for traffic blocked by Google Cloud. A higher priority firewall rule may restrict outbound access. Internet access is allowed if no other firewall rules deny outbound traffic and if the instance has an external IP address or uses a Cloud NAT instance. For more information, see Internet access requirements.

If IPv6 is enabled, the VPC network also has these two implied rules:

- **Implied IPv6 allow egress rule.** An egress rule whose action is allow, destination is ::/0, and priority is the lowest possible (65535) lets any instance send traffic to any destination, except for traffic blocked by Google Cloud. A higher priority firewall rule may restrict outbound access. Internet access is allowed if no other firewall rules deny outbound traffic and if the instance has an external IP address.
- **Implied IPv6 deny ingress rule.** An ingress rule whose action is deny, source is ::/0, and priority is the lowest possible (65535) protects all instances by blocking incoming connections to them. A higher priority rule might allow incoming access.

The implied rules cannot be removed, but they have the lowest possible priorities. You can create rules that override them as long as your rules have higher priorities (priority numbers less than 65535). Because deny rules take precedence over allow rules of the same priority, an ingress allow rule with a priority of 65535 never takes effect.

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

**NEW QUESTION 155**

- (Topic 2)

A large travel company has thus far invested heavily in their technology team. There is strategic pressure on the company to focus on their core business and innovate to survive in certain geographies and thrive in others. They are evaluating whether a move to Google Cloud will be good for them. Which of these reasons would be relevant for them? (choose two answers)

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

**Answer: BC**

**NEW QUESTION 160**

- (Topic 2)

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

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

**Answer: AB**

**Explanation:**

Autoscaling is supported in both Flexible and Standard environments. Flexible Environment does not 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 164**

- (Topic 2)

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 169**

- (Topic 2)

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

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

**Answer:** D

**Explanation:**

"VM instances that only have internal IP addresses (no external IP addresses) can use Private Google Access. They can reach the external IP addresses of Google APIs and services. If you disable Private Google Access, the VM instances can no longer reach Google APIs and services; they can only send traffic within the VPC network."

## Private Google Access

[Send feedback](#)

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

Private Google Access has no effect on instances that have external IP addresses. Instances with external IP addresses can access the internet, according to the [internet access requirements](#). They don't need any special configuration to send requests to the external IP addresses of Google APIs and services.

You enable Private Google Access on a subnet by subnet basis; it's a setting for subnets in a VPC network. To enable a subnet for Private Google Access and to view the requirements, see [Configuring Private Google Access](#).

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

**NEW QUESTION 174**

- (Topic 2)

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 178**

- (Topic 2)

Considering Different Storage and database options e.g. Cloud Datastore, Cloud SQL, Cloud Storage, etc. Which of the following statements is/are correct? (Select two answers)

- A. Cloud DataStore and Cloud SQL have Terabytes + and Terabytes Capacity respectively.
- B. Cloud Bigtable and Cloud Storage both have Petabytes + capacity.
- C. Cloud Bigtable and Cloud Storage both have not Petabytes + capacity.
- D. None of the above.

**Answer:** AB

**NEW QUESTION 179**

- (Topic 2)

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

- A. Circular Scaling takes a round-robin approach to allocate and destroy VMs.

- B. Triangular Scaling takes an automated average of Cost, Effort, and Time.
- C. Vertical Scaling
- D. Horizontal Scaling

**Answer:** D

**Explanation:**

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

**NEW QUESTION 180**

- (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 183**

- (Topic 2)

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 185**

- (Topic 2)

Your company has signed up with a cloud provider and you will be using storage and virtual machines with the provider. The provider has provided your organization some expectations for what the service should perform at. What type of agreement provides a guarantee 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 190**

- (Topic 2)

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

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

**Answer:** C

**Explanation:**

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

understanding of machine learning.  
<https://cloud.google.com/bigquery-ml/docs/>

#### NEW QUESTION 191

- (Topic 2)

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

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 193

- (Topic 2)

When you update the function in firebase by deploying updated code, instances for older versions are cleaned up along with build artifacts in and replaced by new instances.

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 194

- (Topic 2)

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

- A. A hybrid version consisting of a runtime plane installed on-premises or in a cloud provider of your choice, and a management plane running in Apigee's cloud.
- B. In this model, API traffic and data are confined within your own enterprise-approved boundaries.
- C. A hosted SaaS version in which Apigee maintains the environment, allowing you to concentrate on building your services and defining the APIs to those services.
- D. There are two types of Flavours in Apigee i.
- E. Apigee & Apigee Hybrid.
- F. All of the above are correct.

**Answer:** D

#### Explanation:

Flavors of Apigee

Apigee comes in the following flavors:

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

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

#### NEW QUESTION 197

- (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 198

- (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 199**

- (Topic 2)

How does a least privilege resource access model contribute to cloud security?

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

**Answer:** D

**Explanation:**

This is the definition of a least privilege model.

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

**NEW QUESTION 203**

- (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 happening quickly thereafter from these new accounts. Different parts of the application are being explored, all of which are putting a heavy load on the application. What could be a problem and how can you solve it?

- A. A hacker group has hired a bunch of people to create accounts and manually use the system
- B. Use Cloud Asset Inventory to see if there have been changes in the inventory
- 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 restrict the users to known users.
- G. Automated testing tools might still be running and creating accounts
- 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 208**

- (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 212**

- (Topic 2)

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

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

**Answer:** AD

**Explanation:**

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

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

Reference link- [https://en.wikipedia.org/wiki/High\\_availability](https://en.wikipedia.org/wiki/High_availability)

**NEW QUESTION 213**

- (Topic 2)

A client is currently running software on their on-premise systems that is bound by a certain type of license. They are allowed to run the software on virtualized machines. However, they cannot run them on virtualized machines that are shared by two different companies, teams, or projects. What option do they have on Google Cloud?

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

**Answer:** D

**Explanation:**

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

<https://cloud.google.com/compute/docs/nodes/sole-tenant-nodes>

**NEW QUESTION 218**

- (Topic 2)

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

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

**Answer:** CDE

**Explanation:**

Sustained, Committed and Preemptible

vCPU and memory usage for each of these categories can receive discounts VM vCPU and memory usage for each of these categories can receive discounts Sustained-use discounts—Google offers up to 30% off for workloads that run for most of the billing month on GCP services.

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

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

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

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

**NEW QUESTION 219**

- (Topic 2)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 221**

- (Topic 2)

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

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

**Answer:** B

**Explanation:**

Cloud SQL for MySQL:

Features

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

**NEW QUESTION 226**

- (Topic 2)

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

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

**Answer:** B

**Explanation:**

Cloud Bigtable

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

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

**NEW QUESTION 230**

- (Topic 2)

"With cloud messaging you can Customize and deliver messages accordingly to the predetermined time in the user's local time zone." Comment on the above statement.

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

**Answer:** D

**Explanation:**

Firebase Cloud Messaging:

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

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

Key capabilities of Firebase Cloud Messaging:

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

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

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

**NEW QUESTION 232**

- (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 237**

- (Topic 2)

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

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

**Answer: B**

**Explanation:**

Community Cloud – The cloud infrastructure is planned for selective use by a particular community of consumers from organizations that have mutual interests like security needs, policy, and compliance considerations.

Reference link- [https://csrc.nist.gov/glossary/term/community\\_cloud](https://csrc.nist.gov/glossary/term/community_cloud)

**NEW QUESTION 241**

- (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 243**

- (Topic 2)

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 245**

- (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 249**

- (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 254**

- (Topic 1)

The operating systems of some of your organization's virtual machines may have a security vulnerability. How can your organization most effectively identify all virtual machines that do not have the latest security update?

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

**Answer:** A

**Explanation:**

Security Health Analytics and Web Security Scanner detectors generate vulnerabilities findings that are available in Security Command Center. Your ability to view and edit findings is determined by the Identity and Access Management (IAM) roles and permissions you are assigned. For more information about IAM roles in Security Command Center.

Reference link:-

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

**NEW QUESTION 257**

- (Topic 1)

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

Where should your organization locate this virtual machines?

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

**Answer: B**

**Explanation:**

Multi zone is also redundant within the region and it provides the lowest latency.  
Reference link:-  
<https://cloud.google.com/solutions/best-practices-compute-engine-region-selection>

**NEW QUESTION 261**

- (Topic 1)

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 264**

- (Topic 1)

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

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

**Answer: A**

**Explanation:**

Reference: <https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets>  
Graphical user interface, text, application, email Description automatically generated  
<https://cloud.google.com/architecture/migration-to-google-cloud-transferring-your-large-datasets>

**NEW QUESTION 265**

- (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 as required.
- B. Store the information in Secret Manager and give IAM read permissions as re-quired.
- C. Store the information in Kubernetes Secrets and only grant read permissions to users as required.
- D. Encrypt the information and store it in Cloud Storage for centralized 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 269**

- (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 270

- (Topic 1)

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

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

**Answer:** C

**Explanation:**

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

#### NEW QUESTION 271

- (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 272

- (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 Storage
- 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.

# Introduction to the Organization Policy Service

[Send feedback](#)

The Organization Policy Service gives you centralized and programmatic control over your organization's cloud resources. As the [organization policy administrator](#), you will be able to configure constraints across your entire [resource hierarchy](#).

## Benefits

- Centralize control to configure restrictions on how your organization's resources can be used.
- Define and establish guardrails for your development teams to stay within compliance boundaries.
- Help project owners and their teams move quickly without worry of breaking compliance.

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 276

- (Topic 1)

Your organization needs to plan its cloud infrastructure expenditures. Which should your organization do?

- A. Review cloud resource costs frequently, because costs change often based on use
- B. Review cloud resource costs annually as part of planning your organization's overall budget
- C. If your organization uses only cloud resources, infrastructure costs are no longer part of your overall budget
- D. Involve fewer people in cloud resource planning than your organization did for on-premises resource planning

**Answer:** A

#### Explanation:

Review cloud resource costs frequently, because costs change often based on use because One need to know current usage/ trend for planning; While public cloud eliminates capex, and gets into pay as you go model, the usage pattern determines the cloud cost and hence needs to be measured frequently to enable better forecast

### NEW QUESTION 280

- (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:

# Anthos

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

[Try it free](#)[Contact sales](#)

- ✓ Build, deploy, and optimize apps on GKE and VMs anywhere—simply, flexibly, and securely ✖ Rectangular Snip
- ✓ Consistent development and operations experience for hybrid and multicloud environments
- ✓ Achieve up to 4.8x ROI within 3 years according to the [Forrester Total Economic Impact study](#)
- ✓ Accelerate your VM-based app [migration journey](#) to containers

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

## NEW QUESTION 281

- (Topic 1)

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

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

**Answer:** B

### Explanation:

An identity platform is a modern solution for managing the identities of users and devices in a centralized fashion.

Reference: <https://www.okta.com/blog/2021/07/what-is-an-identity-platform/#:~:text=An%20identity%20platform%20is%20a,%2C%20integrations%2C%20and%20platform%20services>

## NEW QUESTION 285

- (Topic 1)

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

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

**Answer:** C

### Explanation:

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

## NEW QUESTION 289

- (Topic 1)

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

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

**Answer:** D

**Explanation:**

Features of Cloud Spanner

Reference: <https://k21academy.com/google-cloud/cloud-sql-vs-cloud-spanner/>

**NEW QUESTION 290**

- (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 292**

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual Cloud-Digital-Leader Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the Cloud-Digital-Leader Product From:

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

## Money Back Guarantee

### Cloud-Digital-Leader Practice Exam Features:

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