Google

GCP-CDL
Google Cloud Digital Leader Certification Exam

For More Information – Visit link below:
https://www.examsempire.com/

Product Version

1. **Up to Date products, reliable and verified.**
2. **Questions and Answers in PDF Format.**

https://examsempire.com/

Latest Version: 6.0

Question: 1

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

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

Answer: B

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

SecureAuth is dedicated to providing the **industry-leading** enhanced **support** ensuring the long term success of your SecureAuth SaaS IAM deployment.

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

Question: 2

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

A. VPN tunnels
B. Identity Platform
C. Compute Engine firewall rules

D. Private Google Access

Answer: B

Explanation:
An identity platform is a modern solution for managing the identities of users and devices in a centralized fashion.
Reference: https://www.okta.com/blog/2021/07/what-is-an-identityplatform/#:~:text=An%20identity%20platform%20is%20a,%2C%20integrations%2C%20and%20platform%20services

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

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

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

Question: 3

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

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

Answer: A

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

What is BigQuery ML?

Send feedback

BigQuery ML lets you create and execute machine learning models in BigQuery using standard SQL queries. BigQuery ML democratizes machine learning by letting SQL practitioners build models using existing SQL tools and skills. BigQuery ML increases development speed by eliminating the need to move data.

BigQuery ML functionality is available by using:

- The Google Cloud console
- The `bq` command-line tool
- The BigQuery REST API
- An external tool such as a Jupyter notebook or business intelligence platform

Machine learning on large datasets requires extensive programming and knowledge of ML frameworks. These requirements restrict solution development to a very small set of people within each company, and they exclude data analysts who understand the data but have limited machine learning knowledge and programming expertise.

BigQuery ML empowers data analysts to use machine learning through existing SQL tools and skills. Analysts can use BigQuery ML to build and evaluate ML models in BigQuery. Analysts don't need to export small amounts of data to spreadsheets or other applications or wait for limited resources from a data science team.

Question: 4

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

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

Answer: C

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

Question: 5

Your organization is developing a plan for migrating to Google Cloud.
What is a best practice when initially configuring your Google Cloud environment?

A. Create a project via Google Cloud Console per department in your company

B. Define your resource hierarchy with an organization node on top
C. Create projects based on tem members' requests
D. Make every member of your company the project owner

Answer: B

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

Question: 6

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

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.

Question: 7

How should a multinational organization that is migrating to Google Cloud consider security and privacy regulations to ensure that it is in compliance with global standards?

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

Answer: A

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

Question: 8

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

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

D. Build a homogenous infrastructure at each subsidiary, and invest in hiring more engineers.

Answer: B

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

Question: 9

What is the difference between Standard and Coldline storage?

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

Answer: D

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

The main characteristics of Coldline are as follows:

- SLA guarantees 99% data availability (of ten thousand hours, data can be offline for up to 100 hours).
- Monthly fee per GB stored is between $0.004-$.0.014, depending on a region.
- The minimum storing period is 90 days. If you delete data earlier, you have to pay for the remaining time. For example, if you had uploaded 100GB of data and then deleted it after 30 days, you need to pay extra $1.4 as an early deletion fee ($0.004 (or $0.014) * 100 GB * 2 months).
- Data retrieval fee is mandatory and costs $0.05 per GB.

Its low price and a long minimum storing period make Nearline the best solution for data that is unlikely to be accessed more than once a year, if ever:

- Data archive.
- Disaster recovery storage.
- Outdated backups storage.

When using Coldline or Nearline Storage, you also pay more for requests, e.g., retrieval of metadata or download commands.

Question: 10

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

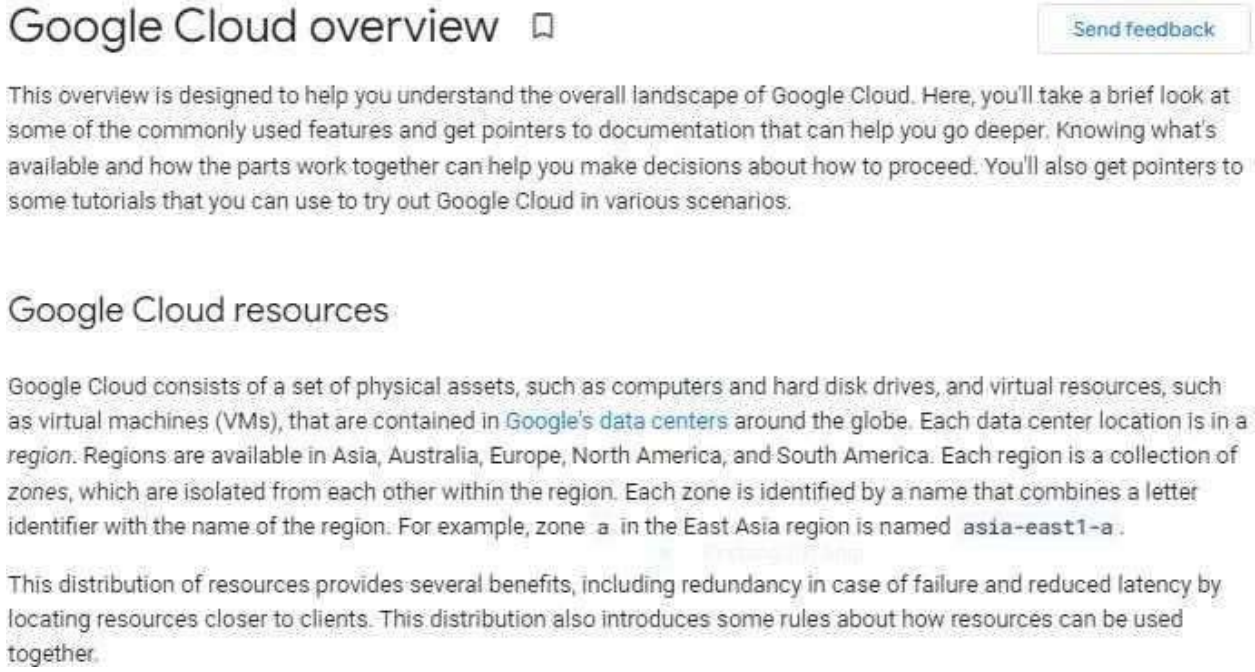

Google Cloud overview

Send feedback

This overview is designed to help you understand the overall landscape of Google Cloud. Here, you'll take a brief look at some of the commonly used features and get pointers to documentation that can help you go deeper. Knowing what's available and how the parts work together can help you make decisions about how to proceed. You'll also get pointers to some tutorials that you can use to try out Google Cloud in various scenarios.

Google Cloud resources

Google Cloud consists of a set of physical assets, such as computers and hard disk drives, and virtual resources, such as virtual machines (VMs), that are contained in Google's data centers around the globe. Each data center location is in a *region*. Regions are available in Asia, Australia, Europe, North America, and South America. Each region is a collection of *zones*, which are isolated from each other within the region. Each zone is identified by a name that combines a letter identifier with the name of the region. For example, zone `a` in the East Asia region is named `asia-east1-a`.

This distribution of resources provides several benefits, including redundancy in case of failure and reduced latency by locating resources closer to clients. This distribution also introduces some rules about how resources can be used together.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

Check our Customer Testimonials and ratings available on every product page.

Visit our website.

https://examsempire.com/