



Google

Exam Questions Cloud-Digital-Leader

Google Cloud Digital Leader exam

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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 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 3

- (Topic 1)

Your organization needs to allow a production job to have access to a BigQuery dataset. The production job is running on a Compute Engine instance that is part of an instance group.

What should be included in the IAM Policy on the BigQuery dataset?

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

Answer: C

Explanation:

When an identity calls a Google Cloud API, BigQuery requires that the identity has the appropriate permissions to use the resource. You can grant permissions by granting roles to a user, a group, or a service account.

Reference link- <https://cloud.google.com/bigquery/docs/access-control>

NEW QUESTION 4

- (Topic 1)

Your organization runs all its workloads on Compute Engine virtual machine instances. Your organization has a security requirement: the virtual machines are not allowed to access the public internet. The workloads running on those virtual machines need to access BigQuery and Cloud Storage, using their publicly accessible interfaces, without violating the security requirement.

Which Google Cloud product or feature should your organization use?

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

Answer: D

Explanation:

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

Configuring Private Google Access

[Send feedback](#)

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

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

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

See [Private Access Options for Services](#) for background information about Private Google Access and other private connectivity options offered by Google Cloud.

Specifications

A VM interface can send packets to the external IP addresses of Google APIs and services using Private Google Access if all these conditions are met:



- The VM interface is connected to a subnet where Private Google Access is enabled.
- The VPC network that contains the subnet meets the [network requirements for Google APIs and services](#).
- The VM interface does not have an external IP address assigned.
- The source IP address of packets sent from the VM matches one of the following IP addresses.

If you're sending packets to the [default domains](#):

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

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

NEW QUESTION 5

- (Topic 1)

Your organization needs to minimize how much it pays for data traffic from the Google network to the internet. What should your organization do?

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

Answer: A

Explanation:

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

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

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

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

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

NEW QUESTION 6

- (Topic 1)

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

How would you restrict all virtual machines from having an external IP address?

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

Answer: A

Explanation:

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

NEW QUESTION 7

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 8

- (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 9

- (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 10

- (Topic 1)

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

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

Answer: D

Explanation:

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

- * 1. Gain centralized visibility and control
- * 2. Discover misconfigurations and vulnerabilities
- * 3. Report on and maintain compliance
- * 4. Detect threats targeting your Google Cloud assets <https://cloud.google.com/security-command-center>

NEW QUESTION 10

- (Topic 1)

Your organization needs to establish private network connectivity between its on-premises network and its workloads running in Google Cloud. You need to be able to set up the connection as soon as possible.

Which Google Cloud product or feature should you use?

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

Answer: A

Explanation:

Private Google Access for on-premises hosts provides a way for on-premises systems to connect to Google APIs and services by routing traffic through a Cloud VPN tunnel. Reference: <https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid>

NEW QUESTION 12

- (Topic 1)

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

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

Answer: A

Explanation:

Explanation

=> Cloud Pub/Sub- Cloud Pub/Sub is the best to be the end-point for ingesting large amounts of data. It will grow as required, can stream data to downstream systems, and can also work with intermittently available backends.

=> Cloud Dataflow- supports streaming data and therefore is an appropriate option for processing the data that is ingested.

=> BigQuery- BigQuery also supports streaming data and its possible to do real time analytics on it.

=> DataStudio- DataStudio and Looker are for visualization. They don't have any in-built analysis.

=> Cloud Functions- Cloud Functions is a useful serverless endpoint. However, Pub/Sub is better in this case because it can also retain messages for a set period if it was not possible to deliver it first time.

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

NEW QUESTION 16

- (Topic 1)

What is the difference between Standard and Coldline storage?

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

Answer: D

Explanation:

Reference: <https://www.msp360.com/resources/blog/google-cloud-nearline-storage-vs-coldline-vs-standard/>

Google Cloud Coldline is a new cold-tier storage for archival data with access frequency of less than once per year. Unlike other cold storage options, Nearline has no delays prior to data access, so now it is the leading solution among competitors.

NEW QUESTION 20

- (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 24

- (Topic 1)

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.

NEW QUESTION 29

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 34

- (Topic 1)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 39

- (Topic 1)

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

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

Answer: D

Explanation:

Bare Metal Solution for Oracle

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

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

NEW QUESTION 40

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 44

- (Topic 1)

Your large and frequently changing organization's user information is stored in an on-premises LDAP database. The database includes user passwords and group and organization membership.

How should your organization provision Google accounts and groups to access Google Cloud resources?

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

Answer: C

Explanation:

You can run a single instance of Google Cloud Directory Sync to synchronize user accounts and groups to Google Cloud.

Reference: <https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-introduction> Text Description automatically generated <https://support.google.com/a/answer/106368?hl=en>

NEW QUESTION 49

- (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 53

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 58

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 60

- (Topic 1)

What would provide near-unlimited availability of computing resources without requiring your organization to procure and provision new equipment?

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

Answer: A

Explanation:

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

NEW QUESTION 65

- (Topic 1)

Your organization needs to process large amounts of data from an online application that operates continuously. You do not want to be required to provision infrastructure or create server clusters. What should your organization choose?

- A. Compute Engine with BigQuery
- B. Dataproc
- C. Google Kubernetes Engine with Cloud Bigtable
- D. Dataflow

Answer: D

Explanation:

You do not want to be required to provision infrastructure or create server clusters. Because Unified stream and batch data processing that's serverless, fast, and cost-effective.

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

NEW QUESTION 68

- (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 72

- (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 76

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 77

- (Topic 1)

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

- A. Use Private Service Access within Google Cloud.

- B. Use VPC Peering to share resources privately between your two organizations.
- C. Use public IP addresses as before.
- D. It will automatically be routed internally only.
- E. Use VPC Shared Networks to share common resources.

Answer: B

Explanation:

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

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

-> Private Google Access is only to access Google APIs and services

References:

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

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

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

NEW QUESTION 79

- (Topic 3)

An organization is looking for a storage solution that will help them serve content to users worldwide. They need a solution that offers a high level of availability. What feature of Cloud Storage would they benefit from?

- A. Global metadata
- B. Object versioning
- C. Data encryption
- D. Multi-regional storage

Answer: D

NEW QUESTION 84

- (Topic 3)

An organization wants to collect metrics and metadata from their cloud applications and put them into dashboards.

Which Google Cloud tool should they use?

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

Answer: A

Explanation:

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

NEW QUESTION 89

- (Topic 3)

What is monitoring within the context of cloud operations?

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

Answer: B

NEW QUESTION 93

- (Topic 3)

An organization needs to search an application's source code to identify a potential issue. The application is distributed across multiple containers.

Which Google Cloud product should the organization use?

- A. Google Cloud Console
- B. Cloud Trace
- C. Cloud Monitoring
- D. Cloud Logging

Answer: B

Explanation:

Cloud Trace is supposed to be the correct answer. It's an application performance management tool. It's a Google solution for monitoring application performance. It is a distributed tracing system that helps developers debug or fix and optimize their code

NEW QUESTION 95

- (Topic 3)

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

- A. Number of available regions
- B. Percentage of uptime

- C. Speed of response time
- D. Number of downtime incidents

Answer: B

NEW QUESTION 97

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 101

- (Topic 3)

Why is data stored in Google Cloud secure and private?

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

Answer: C

Explanation:

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

NEW QUESTION 104

- (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 106

- (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 110

- (Topic 3)

An organization wants to use BigQuery data analytics to understand their website performance, but wants to move only some data into the cloud. Which environment should the organization use?

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

Answer: D

Explanation:

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

NEW QUESTION 111

- (Topic 3)

After rolling out a new update, an organization found a minor bug in its online video game. How should the organization approach this bug while following SRE principles?

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

Answer: A

Explanation:

<https://www.blameless.com/sre/sre-principles>

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

NEW QUESTION 112

- (Topic 3)

An organization is struggling to keep up with the growth of their application which is running on legacy infrastructure. What might be holding them back?

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

Answer: D

Explanation:

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

NEW QUESTION 115

- (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 120

- (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 125

- (Topic 3)

How does a large hotel chain benefit from storing their customer reservation data in the cloud?

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

Answer: C

NEW QUESTION 128

- (Topic 3)

An organization is planning its cloud expenditure. What should the organization do to control costs?

- A. Consider cloud resource costs as capital expenditure in annual planning.
- B. Use only cloud resources; they have no cloud infrastructure costs.
- C. Review cloud resource costs frequently because costs depend on usage.
- D. Assess cloud resources costs only when SLO is not met by their cloud provider.

Answer: C

NEW QUESTION 129

- (Topic 3)

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

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

Answer: B

NEW QUESTION 131

- (Topic 3)

An organization operates their entire IT infrastructure from Google Cloud. What should they do to prepare for data breaches?

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

Answer: C

NEW QUESTION 134

- (Topic 3)

A global organization is developing an application to manage payments and online bank accounts in multiple regions. Each transaction must be handled consistently in their database, and they anticipate almost unlimited growth in the amount of data stored. Which Google Cloud product should the organization choose?

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

Answer: B

NEW QUESTION 137

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 140

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 143

- (Topic 3)

A retail company stores their product inventory in a legacy system. Often, customers find products on the company's website and want to purchase them in-store. However, when they arrive, they discover that the products are out of stock.

How could the company benefit from using an application programming interface (API)?

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

Answer: D

Explanation:

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

NEW QUESTION 147

- (Topic 3)

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

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

Answer: A

NEW QUESTION 151

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 153

- (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 156

- (Topic 3)

An organization has an on-premises IT infrastructure. Their customer-facing application repeatedly fails during peak usage. What could be causing this issue?

- A. A serverless compute function struggles to scale.
- B. The application contains unclean data.
- C. They don't have enough servers to meet the demand.
- D. The application is only configurable on-premises.

Answer: C

NEW QUESTION 160

- (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 163

- (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 166

- (Topic 3)

An organization wants full control of their virtual machine infrastructure for a custom home-grown application with a product that autoscales and automatically updates.

Which Google Cloud product or solution should the organization use?

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

Answer: C

Explanation:

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

NEW QUESTION 168

- (Topic 3)

A food delivery service needs access to real-time menu information from all partner restaurants. They also need to share customer order information with the restaurants in real time.

What should the organization use?

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

Answer: B

NEW QUESTION 169

- (Topic 3)

An organization is migrating their business applications from on-premises to the cloud. How could this impact their operations and personnel costs?

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

Answer: A

NEW QUESTION 170

- (Topic 3)

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

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

Answer: B

NEW QUESTION 173

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 178

- (Topic 3)

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

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

Answer: B

NEW QUESTION 182

- (Topic 3)

When an organization adopts cloud technology, how does their total cost of ownership (TCO) shift?

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

Answer: C

NEW QUESTION 183

- (Topic 3)

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

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

Answer: A

Explanation:

Virtual machines - you can install customized OS Containers - about applications

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

NEW QUESTION 185

- (Topic 3)

What is logging within the context of cloud technology?

- A. Writing application and operating system events as text
- B. Monitoring network and resource limitations
- C. Tracking source code across an organization
- D. Recording infrastructure and hardware expenditure

Answer: A

Explanation:

Cloud Logging is a fully managed service that allows you to store, search, analyze, monitor, and alert on logging data and events from Google Cloud and Amazon Web Services

NEW QUESTION 187

- (Topic 3)

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

- A. Customer expectations can be adjusted without using marketing tools
- B. Seamless changes can be made without causing any application downtime.
- C. GKE handles version control seamlessly and out of the box
- D. GKE is well suited for all monolithic applications

Answer: B

Explanation:

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

NEW QUESTION 191

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 192

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 196

- (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 200

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 202

- (Topic 3)

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

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

Answer: A

NEW QUESTION 206

- (Topic 3)

An organization is using machine learning to make predictions. One of their datasets mistakenly includes mislabeled data. How will the prediction be impacted?

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

Answer: B

NEW QUESTION 208

- (Topic 3)

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

- A. Maintaining overall system operability
- B. Maintaining customer-facing content
- C. Monitoring data center servers
- D. Monitoring computer networks

Answer: B

NEW QUESTION 211

- (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 216

- (Topic 3)

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

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

Answer: B

NEW QUESTION 218

- (Topic 3)

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

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

Answer: C

NEW QUESTION 223

- (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 224

- (Topic 3)

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

- A. Ransomware
- B. Distributed Denial of Service
- C. Spamming

D. Phishing

Answer: D

Explanation:

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

NEW QUESTION 229

- (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 230

- (Topic 2)

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

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

Answer: AB

Explanation:

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

NEW QUESTION 234

- (Topic 2)

Your client is building a custom machine learning pipeline to identify lesions in the lungs based on x-rays. Different teams of data scientists are sharing common source data and building many ver-sions of ML models. Which of these Cloud Storage options would be best for them?

- A. Retain the data in use in a single region bucket with nearline storag
- B. Retain the data in use in a dual-region bucket.
- C. Retain the data in use in a single region bucket with standard storage.
- D. Retain the data in use in a multi-region bucket.
- E. Retain the data in use in a dual-region bucket.

Answer: B

Explanation:

Integrated repository for analytics and ML: The highest level of availability and performance within a single region is ideal for compute, analytics, and machine learning workloads in a particular region. Cloud Storage is also strongly consistent, giving you confidence and accuracy in analytics workloads.

Standard storage

Standard storage is best for data that is frequently accessed ("hot" data) and/or stored for only brief periods of time.

When used in a region, Standard storage is appropriate for storing data in the same location as [Google Kubernetes Engine clusters](#) or [Compute Engine instances](#) that use the data. Co-locating your resources maximizes the performance for data-intensive computations and can reduce network charges.

When used in a dual-region, you still get optimized performance when accessing Google Cloud products that are located in one of the associated regions, but you also get the improved availability that comes from storing data in geographically separate locations.

When used in a multi-region, Standard storage is appropriate for storing data that is accessed around the world, such as serving website content, streaming videos, executing interactive workloads, or serving data supporting mobile and gaming applications.

Availability

The availability of Standard storage data is:

Location Type	Availability SLA ¹	Typical monthly availability
multi-region	99.95%	>99.99%
dual-region	99.95%	>99.99%

Table Description automatically generated with medium confidence <https://cloud.google.com/storage/docs/storage-classes>

NEW QUESTION 238

- (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 answer)

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

Answer: BC

NEW QUESTION 239

- (Topic 2)

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

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

Answer: AB

Explanation:

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

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

NEW QUESTION 241

- (Topic 2)

You are working with a government agency. A web application serves users of the country. It allows citizens to receive certain services in providing their national identity. Citizens have complained that they are seeing delays in web page loading compared to before. On investigating, they are seeing a lot of spurious traffic coming in from a few IPs which they have identified as foreign. What should they do?

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

Answer: B

Explanation:

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

NEW QUESTION 243

- (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 247

- (Topic 2)

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 251

- (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 256

- (Topic 2)

You are storing sensitive information in a Cloud Storage bucket. For legal reasons, you need to be able to record all requests that read any of the stored data. You want to make sure you comply with these requirements. What should you do?

- A. Scan the bucket using the Data Loss Prevention API.

- B. Enable Data Access audit logs for the Cloud Storage API.
- C. Enable the Identity Aware Proxy API on the project.
- D. Allow only a single Service Account access to read the data.

Answer: B

Explanation:

Logged information

Your Google Cloud projects contain only the audit logs for resources that are directly within the Cloud project. Other Google Cloud resources, such as folders, organizations, and billing accounts, contain the audit logs for the entity itself.

Available audit logs

The following types of audit logs are available for Cloud Storage:

- **Admin Activity audit logs:** Entries for `ADMIN_WRITE` operations that modify the configuration or metadata of a Cloud project, bucket, or object. You can't disable Admin Activity audit logs.
- **Data Access audit logs:** Entries for operations that modify objects or read a Cloud project, bucket, or object. There are several sub-types of Data Access audit logs:
 - `ADMIN_READ`: Entries for operations that read the configuration or metadata of a Cloud project, bucket, or object.
 - `DATA_READ`: Entries for operations that read an object.
 - `DATA_WRITE`: Entries for operations that create or modify an object.

To receive Data Access audit logs, you must **explicitly enable** them.

For fuller descriptions of the audit log types, see [Types of audit logs](#).

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

NEW QUESTION 257

- (Topic 2)

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

- A. About 12 days.
- B. About 2 years.
- C. About 100 days.
- D. About 24 hours.

Answer: C

Explanation:

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

Total data transferred (in bits) = $100 * 1,000,000,000,000 * 8$ bits

Speed = 100Mbps = $100 * 1,000,000$. i.e. 100 million bits are transferred per second. Hence time taken to transfer all the data = $\text{Total Data}/\text{Speed} = 8,000,000$ seconds.

Number of seconds in a day = $24 * 60 * 60 = 86,400$

Total time taken in days = $8,000,000 / 86,400 = 92.59$ days

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

NEW QUESTION 261

- (Topic 2)

Google Cloud Platform (GCP) provides three main compliance resource webpages. What are they? (Select Three Answer)

- A. Compliance Reports Manager
- B. Support Hub
- C. Compliance Offerings
- D. GDPR Home Page
- E. TechCentral

Answer: ACD

Explanation:

Compliance Reports Manager, GDPR Home Page, Compliance Offerings GCP provides three main compliance resource webpages

Compliance Reports Manager

– <https://cloud.google.com/security/compliance/compliance-reports-manager>

Compliance Reports Manager

Google Cloud's industry-leading security, third-party audits and certifications, documentation, and contract commitments help support your compliance. Compliance reports manager provides you with easy, on-demand access to these critical compliance resources, at no additional cost. Key resources include our latest ISO/IEC certificates, SOC reports, and self assessments.

Select resources may require sign-in with your Google Cloud or Google Workspace account. If you would like to access previous reports please reach out to [support](#) for more information. Anything marked "Google Confidential Information" is shared subject to the confidentiality obligations described in the customer or partner agreement(s) covering Cloud Services. Please contact your sales representative for permission to share confidential resources outside of your organization with customers or other third parties not expressly permitted by your agreement.

Text, timeline Description automatically generated
Compliance Offerings – <https://cloud.google.com/security/compliance/offerings>

Compliance offerings

To help you with compliance and reporting, we share information, best practices, and easy access to documentation.

Our products regularly undergo independent verification of security, privacy, and compliance controls, achieving certifications against global standards to earn your trust. We're constantly working to expand our coverage.

This site contains information about Google's certifications and compliance standards it satisfies as well as general information about certain region or sector-specific regulations.

Text Description automatically generated
GDPR Resource Center – <https://cloud.google.com/security/gdpr/resource-center> At Google Cloud, we champion initiatives that prioritize and improve the security and privacy of customer personal data, and want you, as a Google Cloud customer, to feel confident using our services in light of GDPR requirements. If you partner with Google Cloud, we will support your GDPR compliance efforts

NEW QUESTION 265

- (Topic 2)

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

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

Answer: C

Explanation:

Firestore Crashlytics:

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

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

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

NEW QUESTION 267

- (Topic 2)

Which of the followings are core components of Anthos?

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

Answer: D

Explanation:

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

NEW QUESTION 272

- (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 273

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 277

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 279

- (Topic 2)

Your client has an on-premises data center. Due to technical limitations, they are unable to scale globally. They have decided to adopt the public cloud. However, they don't want to be locked into any one vendor and, therefore, would like to work with multiple cloud providers. They have used open source container technologies and would like to continue using them.

- A. Cloud Run which supports containers and can scale in a serverless fashion
- B. Kubernetes that runs containers as their core workloads
- C. AppEngine Flexible Environment which supports containers
- D. Anthos that runs containers as their core workloads

Answer: D

Explanation:

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

Anthos enables you to manage GKE clusters and workloads running on virtual machines across environments. You get consistent managed Kubernetes experience with simple installs as well as upgrades validated by Google. Anthos can run on your existing virtualized infrastructure and [bare metal](#) servers without a hypervisor layer. Anthos simplifies your application stack, reduces the costs associated with licensing a hypervisor, and decreases time spent learning new skills.

NEW QUESTION 281

- (Topic 2)

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

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

Answer: C

Explanation:

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

Storage Transfer Service is a product that enables you to:

- Move or backup data to a Cloud Storage bucket either from other cloud storage providers or from a local or cloud POSIX file system.
- Move data from one Cloud Storage bucket to another, so that it is available to different groups of users or applications.
- Move data from Cloud Storage to a local or cloud file system
- Move data between file systems.
- Periodically move data as part of a data processing pipeline or analytical workflow.

Storage Transfer Service provides options that make data transfers and synchronization easier. For example, you can:

- Schedule one-time transfer operations or recurring transfer operations.
- Delete existing objects in the destination bucket if they don't have a corresponding object in the source.
- Delete data source objects after transferring them.
- Schedule periodic synchronization from a data source to a data sink with advanced filters based on file creation dates, filenames, and the times of day you prefer to import data.

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

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

NEW QUESTION 285

- (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 288

- (Topic 2)

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

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

Answer: A

Explanation:

Service-Level Agreement (SLA)

At Google, we distinguish between an SLO and a Service-Level Agreement (SLA). An SLA normally involves a promise to someone using your service that its availability SLO should meet a certain level over a certain period, and if it fails to do so then some kind of penalty will be paid. This might be a partial refund of the service subscription fee paid by customers for that period, or additional subscription time added for free. The concept is that going out of SLO is going to hurt the service team, so they will push hard to stay within SLO. If you're charging your customers money, you will probably need an SLA.

Because of this, and because of the principle that availability shouldn't be much better than the SLO, the availability SLO in the SLA is normally a looser objective than the internal availability SLO. This might be expressed in availability numbers: for instance, an availability SLO of 99.9% over one month, with an internal availability SLO of 99.95%. Alternatively, the SLA might only specify a subset of the metrics that make up the internal SLO.

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

NEW QUESTION 289

- (Topic 2)

Which of the following statements describe the features of a preemptible VM in-stance? (Select Three Answer)

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

Answer: CDE

Explanation:

Instance is alive for no more than 24 hours, Can be pre-empted with a 30 second notice, Discounted Significantly.

Preemptible VM is an instance that you can create and run at a lower cost than normal instances.

However, Compute Engine might stop (pre-empt) these instances if it requires access to those resources for other tasks. Preemptible instances are excess Compute Engine capacity, so their availability varies with usage.

Live at most 24 hours Can be pre-empted with a 30 second notification via API and are Discounted significantly

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

NEW QUESTION 290

- (Topic 2)

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

- A. Kubernetes uses Docker to deploy, manage, and scale containerized applications.
- B. Difference between Docker and Kubernetes relates to the role each play in con-tainerizing and running your applications
- C. Kubernetes can be used with or without Docker.
- D. All of the above.

Answer: D

Explanation:

Kubernetes vs. Docker

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

Docker lets you put everything you need to run your application into a box that can be stored and opened when and where it is required. Once you start boxing up your applications, you need a way to manage them; and that's what Kubernetes does. Kubernetes is a Greek word meaning 'captain' in English. Like the captain is responsible for the safe journey of the ship in the seas, Kubernetes is responsible for carrying and delivering those boxes safely to locations where they can be used.

- Kubernetes can be used with or without Docker.

- Docker is not an alternative to Kubernetes, so it's less of a "Kubernetes vs. Docker" question. It's about using Kubernetes with Docker to containerize your applications and run them at scale.

- The difference between Docker and Kubernetes relates to the role each play in containerizing and running your applications.

- Docker is an open industry standard for packaging and distributing applications in containers.

- Kubernetes uses Docker to deploy, manage, and scale containerized applications.

NEW QUESTION 291

- (Topic 2)

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

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

Answer: D

Explanation:

Preemptible instances function like normal instances but have the following limitations:

-> Compute Engine might stop preemptible instances at any time due to system events. The probability that Compute Engine will stop a preemptible instance for a system event is generally low, but might vary from day to day and from zone to zone depending on current conditions.

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

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

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

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

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

NEW QUESTION 292

- (Topic 2)

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

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

Answer: A

Explanation:

Graphical user interface, text, application, email Description automatically generated

DevOps measurements for reliability and system health

DevOps teams can track system reliability, quality, and overall health using a few key metrics. In DevOps organizations, site reliability engineers, operations engineers, software developers, project managers, and engineering leadership will all find value in these measurements.

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

NEW QUESTION 295

- (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 296

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 299

- (Topic 2)

Which of the following methods should you use when you require a dynamic way of provisioning VMs on Compute Engine when it is observed that the exact specifications will be in a dedicated configuration file and you want to follow Google's recommended practices.

- A. Managed Instance Group
- B. Deployment Manager
- C. Cloud Composer
- D. Unmanaged Instance Group

Answer: B

Explanation:

The question is about a dynamic way to provision VM, it can be achieved by a Deployment manager or by using terraform. MIG is creating multiple machines based on templates by load balancing

NEW QUESTION 304

- (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 307

- (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 answer)

- 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

NEW QUESTION 309

- (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 312

- (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 315

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 318

- (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. BigQuary
- B. Spanner
- C. Cloud SQL
- D. Bigtable

Answer: B

Explanation:

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

NEW QUESTION 321

- (Topic 2)

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

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

Answer: B

Explanation:

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

Choose a Google Cloud bare metal migration strategy

[Send feedback](#)

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

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

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

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

NEW QUESTION 324

- (Topic 2)

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

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

Answer: D

Explanation:

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

Below are my top 5 characteristics of a DevOps organization.

- Product based teams over component teams. ... 
- Obsession with Automation over preoccupation with manual work. ...
- Evidence-based over gut feel. ...
- Teamwork over individual work. ...
- Fail fast over delayed learning.

NEW QUESTION 328

- (Topic 2)

What according to you are NOT the key capabilities of In-App Messaging?

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

Answer: C

Explanation:

In-App Messaging

Engage active app users with contextual messages.

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

NEW QUESTION 329

- (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 331

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 333

- (Topic 2)

Cloud SQL is a fully-managed relational database service for MySQL, PostgreSQL and SQL servers, 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 scalable so that your business continues to run without disruption.
- C. With DMS (Database Migration Service) it becomes very easy to Migration of Production 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 334

- (Topic 2)

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

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

Answer: B

Explanation:

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

Reference link- https://csrc.nist.gov/glossary/term/community_cloud

NEW QUESTION 339

- (Topic 2)

You are looking for a one stop reference page for GCP support. What Page would you select?

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

Answer: C

Explanation:

Google provides a page that brings together everything needed around support. Its called the Support Hub

Reference link- <https://cloud.google.com/support-hub>

NEW QUESTION 342

- (Topic 2)

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

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

Answer: B

Explanation:

Cloud BigTable

Key features

High throughput at low latency

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

Cluster resizing without downtime

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

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

NEW QUESTION 346

- (Topic 2)

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

- A. Compute Engine with VMs with either Linux or Windows OS.
- B. App Engine Standard
- C. Cloud Functions
- D. Cloud Run

Answer: A

Explanation:

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

NEW QUESTION 349

- (Topic 2)

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

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

Answer: B

Explanation:

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

-> Pub/Sub allows services to communicate asynchronously, with latencies on the order of 100 milliseconds.

-> Pub/Sub is used for streaming analytics and data integration pipelines to ingest and distribute data. It is equally effective as messaging-oriented middleware for service integration or as a queue to parallelize tasks.

-> Pub/Sub enables you to create systems of event producers and consumers, called publishers and subscribers. Publishers communicate with subscribers asynchronously by broadcasting events, rather than by synchronous remote procedure calls (RPCs).

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

NEW QUESTION 350

- (Topic 2)

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

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

Answer: B

Explanation:

To change the Cloud Billing account for a project, you need to be able to move a project from one Cloud Billing account to another. To accomplish this task, you need permissions adequate to unlink the project from the existing Cloud Billing account AND to link the project to the target Cloud Billing account. Roles with adequate permissions to perform this task: Project Owner or Project Billing Manager on the project, AND Billing Account Administrator or Billing Account User for the target Cloud Billing account

interface, text, application, email Description automatically generated

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

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

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

NEW QUESTION 355

- (Topic 2)

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

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

Answer: C

Explanation:

A top-level idea about Apigee API Management and its offered features can help you solve all questions related to Apigee in Cloud Digital Leader Practice Exam. Apigee is a platform for developing and managing APIs. By fronting services with a proxy

layer, Apigee provides an abstraction or facade for your backend service APIs and provides security, rate limiting, quotas, analytics, and more.

Apigee services: The APIs that you use to create, manage, and deploy your API proxies. Apigee runtime: A set of containerized runtime services in a Kubernetes cluster that Google maintains. All API traffic passes through and is processed by these services.

NEW QUESTION 356

- (Topic 1)

Your organization needs to build streaming data pipelines. You don't want to manage the individual servers that do the data processing in the pipelines. Instead, you want a managed service that will automatically scale with the amount of data to be processed.

Which Google Cloud product or feature should your organization choose?

- A. Pub/Sub
- B. Dataflow
- C. Data Catalog
- D. Dataprep by Trifacta

Answer: B

Explanation:

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

NEW QUESTION 359

- (Topic 1)

You are migrating workloads to the cloud. The goal of the migration is to serve customers worldwide as quickly as possible. According to local regulations, certain data is required to be stored in a specific geographic area, and it can be served worldwide. You need to design the architecture and deployment for your workloads.

What should you do?

- A. Select a public cloud provider that is only active in the required geographic area
- B. Select a private cloud provider that globally replicates data storage for fast data access
- C. Select a public cloud provider that guarantees data location in the required geographic area
- D. Select a private cloud provider that is only active in the required geographic area

Answer: C

Explanation:

The goal of the migration is to serve customers worldwide as quickly as possible. According to local regulations, certain data is required to be stored in a specific geographic area, and it can be served worldwide." This characteristic are inherent to the public cloud provider

NEW QUESTION 364

- (Topic 1)

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

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

Answer: D

Explanation:

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

Premium Support overview

[Send feedback](#)

This page explains the features of Premium Support.

Premium Support is a paid support offering designed for enterprises that run mission critical workloads and require fast response times, platform stability, and increased operational efficiencies.

This overview covers the following aspects of Premium Support:

- How you work with a [Technical Account Manager](#) to optimize your Google Cloud operations and Premium Support experience.
- [Features](#) of the offering, including [support case features](#) and [Customer Aware Support](#).
- [Value Add Services](#) that you can purchase to customize your offering.

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

NEW QUESTION 366

- (Topic 1)

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

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

Answer: C

Explanation:

The pricing calculator can be used to give clear estimates of resource usage.

-> Running test loads is as closely indicative as using the pricing calculator.

-> There are no cloud cost references published, either by Google or CNCF. Even if some companies have published such info. It might not apply to you.

Reference link:- <https://cloud.google.com/products/calculator>

NEW QUESTION 368

- (Topic 1)

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

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

Answer: A

Explanation:

Reference: <https://cloud.google.com/compute/docs/import/importing-virtual-disks>
Graphical user interface, text, application, email Description automatically generated

NEW QUESTION 371

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 374

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 378

- (Topic 1)

Your organization consists of many teams. Each team has many Google Cloud projects. Your organization wants to simplify the management of identity and access policies for these projects.

How can you group these projects to meet this goal?

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

Answer: C

Explanation:

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

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

NEW QUESTION 382

- (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 384

- (Topic 1)

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

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

Answer: A

Explanation:

What is a preemptible instance?

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

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

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

NEW QUESTION 385

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 386

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 390

- (Topic 1)

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

What should you do?

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

Answer: D

Explanation:

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

NEW QUESTION 395

- (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 399

- (Topic 1)

Which Google Cloud product or feature makes specific recommendations based on security risks and compliance violations?

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

Answer: B

Explanation:

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

NEW QUESTION 401

- (Topic 1)

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

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

Answer: B

Explanation:

Enabling Requester Pays is useful, for example, if you have a lot of data you want to make available to users, but you don't want to be charged for their access to that data.

Reference link- <https://cloud.google.com/storage/docs/requester-pays>

NEW QUESTION 404

- (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 409

- (Topic 1)

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

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

Answer: D

Explanation:

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

Migrate workloads to Compute Engine with Migrate for Compute Engine. Migrate from Compute Engine to containers with Migrate for Anthos and GKE.

This method makes sense, for instance, in cases where you want to conduct a data-center migration and migrate all workloads into Compute Engine, and only at a second stage selectively modernize suitable workloads to containers.

Graphical user interface, text, application, email Description automatically generated

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

NEW QUESTION 412

.....

Relate Links

100% Pass Your Cloud-Digital-Leader Exam with Examible Prep Materials

<https://www.exambible.com/Cloud-Digital-Leader-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>