

Cloud-Digital-Leader Dumps

Google Cloud Digital Leader exam

<https://www.certleader.com/Cloud-Digital-Leader-dumps.html>



NEW QUESTION 1

- (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 2

- (Topic 1)

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

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

Answer: C

Explanation:

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

References:

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

-> <https://cloud.google.com/ai-platform/data-labeling/docs>

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 is releasing its first publicly available application in Google Cloud. The application is critical to your business and customers and requires a 2-hour SLA.

How should your organization set up support to minimize costs?

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

Answer: B

Explanation:

Reference: <https://www.secureauth.com/enhanced-support-offering/>

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

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)

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 7

- (Topic 1)

Which of the following is/are true about Anthos?

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

Answer: D

Explanation:

Anthos :

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

- Build, deploy, and optimize apps on GKE and VMs anywhere—simply, flexibly, and securely.

- Consistent development and operations experience for hybrid and multi-cloud environments.

Key features:

* 1. Enterprise-grade container orchestration and management service

* 2. Automate policy and security at scale

- * 3. Fully managed service mesh with built-in visibility
- * 4. Modernizing your security for hybrid and multi-cloud deployments

NEW QUESTION 8

- (Topic 1)

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

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

Answer: C

Explanation:

Deploy your application to App Engine using the `gcloud app deploy` command. This command automatically builds a container image by using the Cloud Build service and then deploys that image to the App Engine flexible environment.

Reference: <https://cloud.google.com/appengine/docs/flexible/nodejs/testing-and-deploying-your-app>

NEW QUESTION 9

- (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 10

- (Topic 1)

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

Why would SaaS be the right choice of service model?

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

Answer: B

Explanation:

Benefits of SaaS

The main benefit of SaaS is that it offloads all infrastructure and application management to the SaaS vendor.

Reference: <https://www.ibm.com/cloud/learn/iaas-paas-saas>

NEW QUESTION 10

- (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 12

- (Topic 1)

You are a program manager for a team of developers who are building an event-driven application to allow users to follow one another's activities in the app. Each time a user adds himself as a follower of another user, a write occurs in the real-time database.

The developers will develop a lightweight piece of code that can respond to database writes and generate a notification to let the appropriate users know that they have gained new followers. The code should integrate with other cloud services such as Pub/Sub, Firebase, and Cloud APIs to streamline the orchestration process. The application requires a platform that automatically manages underlying infrastructure and scales to zero when there is no activity.

Which primary compute resource should your developers select, given these requirements?

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

Answer: B

Explanation:

Reference: <https://firebase.google.com/docs/functions/use-cases>

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

NEW QUESTION 15

- (Topic 1)

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

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

Answer: B

Explanation:

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

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

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

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

References:

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

NEW QUESTION 18

- (Topic 1)

As your organization increases its release velocity, the VM-based application upgrades take a long time to perform rolling updates due to OS boot times. You need to make the application deployments faster.

What should your organization do?

- A. Migrate your VMs to the cloud, and add more resources to them
- B. Convert your applications into containers
- C. Increase the resources of your VMs
- D. Automate your upgrade rollouts

Answer: B

NEW QUESTION 19

- (Topic 1)

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

What Google Cloud product or feature should your organization use?

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 21

- (Topic 1)

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

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

Answer: C

Explanation:

Graphical user interface, text, application, Word, email

Direct Peering overview

[Send feedback](#)

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

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

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

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

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

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

NEW QUESTION 22

- (Topic 1)

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

What should your organization do?

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

Answer: C

Explanation:

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

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

Sharing committed use discounts across projects

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

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

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

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

NEW QUESTION 27

- (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 30

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 35

- (Topic 1)

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

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

Answer: AB

Explanation:

Resource hierarchy	Structure and organize your resource hierarchy for fine-grained management and cost allocation using organizations, folders, projects, and labels.
Billing access control	Enforce organizational policies with granular permissions at different levels in the resource hierarchy to control who can spend and who has administrative and cost-viewing permissions.

Description automatically generated with medium confidence

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

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

NEW QUESTION 40

- (Topic 1)

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

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

Answer: C

Explanation:

Explanation

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

Database Migration Service

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

Get started

Migration guide

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

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

NEW QUESTION 43

- (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 48

- (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 52

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 55

- (Topic 1)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 60

- (Topic 1)

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

- A. The model is under fitted

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

Answer: D

Explanation:

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

NEW QUESTION 62

- (Topic 3)

How would an organization benefit from using Looker?

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

Answer: D

Explanation:

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

NEW QUESTION 67

- (Topic 3)

How is privacy defined in the context of cloud technology?

- A. Restrictions on data access and sharing
- B. Procedures to authenticate user identity
- C. Susceptibility to data breaches and cyber attacks
- D. Compliance with regulatory standards

Answer: A

NEW QUESTION 71

- (Topic 3)

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

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

Answer: C

NEW QUESTION 76

- (Topic 3)

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

What benefit could the organization experience?

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

Answer: D

NEW QUESTION 81

- (Topic 3)

An organization wants to build autoscaling web applications without having to manage application infrastructure. Which Google Cloud product should they use?

- A. App Engine
- B. AutoML
- C. Anthos
- D. Apigee

Answer: A

Explanation:

Per Google docs, App Engine, allows for "freeing up your developers with zero server management and zero configuration deployments".

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

NEW QUESTION 86

- (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 87

- (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 89

- (Topic 3)

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

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

Answer: A

Explanation:

<https://cloud.google.com/natural-language>

Use entity analysis to find and label fields within a document—including emails, chat, and social media—and then sentiment analysis to understand customer opinions to find actionable product and UX insights.

NEW QUESTION 90

- (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 91

- (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 95

- (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 99

- (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 100

- (Topic 3)

An organization's public cloud provider failed to meet their SLA of 99.99% availability. What is the potential impact on the organization?

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

Answer: C

NEW QUESTION 102

- (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 105

- (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 109

- (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 112

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 115

- (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 120

- (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 123

- (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 128

- (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 132

- (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 134

- (Topic 3)

Which policy helps Google Cloud keep customer data private?

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

Answer: B

NEW QUESTION 138

- (Topic 3)

An organization is moving away from an on-premises infrastructure. Instead, they want to create, access, and share information virtually in the cloud. What should the organization consider?

- A. Built-in security when moving their data to the cloud
- B. Replacing their perimeter security with data encryption keys
- C. Optimizing cost-management with a capital expenditure model
- D. Increased hardware capacity when moving their data to the cloud

Answer: A

NEW QUESTION 143

- (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 146

- (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 149

- (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 151

- (Topic 3)

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

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

Answer: C

NEW QUESTION 154

- (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 156

- (Topic 3)

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

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

Answer: D

NEW QUESTION 157

- (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 162

- (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 165

- (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 168

- (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 172

- (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 173

- (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 174

- (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 178

- (Topic 3)

What does Cloud Logging help an organization do?

- A. Analyze live source code and log code updates.
- B. Deploy infrastructure as code.
- C. Analyze logs and accelerate application troubleshooting.
- D. Manage storage of custom VM images.

Answer: C

NEW QUESTION 183

- (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 187

- (Topic 3)

An organization wants to create a new application in the cloud to replace an existing on-premises application. Which application modernization approach should the organization use?

- A. Move the application to the cloud, and then change it.
- B. Change their application, and then move it to the cloud.
- C. Invent in greenfield.
- D. Invent in brownfield.

Answer: D

Explanation:

This approach carries over as much custom components as possible from the source system and minimizes initial reengineering efforts.

NEW QUESTION 191

- (Topic 3)

How does Cloud SQL help organizations create business insights?

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

Answer: A

Explanation:

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

NEW QUESTION 196

- (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 197

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 202

- (Topic 2)

The customer has applications that do data processing on-premise. They have been built using Hadoop and Spark. What product should I use on Google Cloud?

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

Answer: A

Explanation:

Because Dataproc is used to run Hadoop/Spark workloads

NEW QUESTION 203

- (Topic 2)

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

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

Answer: B

Explanation:

Because unsecured third-party systems are a cybersecurity threat.

NEW QUESTION 207

- (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 212

- (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 geogra-phies 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 215

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 216

- (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 220

- (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 221

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 223

- (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 = Total Data/Speed = 8,000,000 seconds.

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

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

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

NEW QUESTION 225

- (Topic 2)

DriveSuper Inc. teaches its clients to drive cars and bikes and helps them get their license. They are planning to build a mobile application where users can sign

up, plan their schedules, and take stock of progress. They want the onboarding process to be smooth and frictionless, giving users a great experience from the get-go. They want this done as quickly as possible and not be expensive. What is their best option on Google Cloud?

- A. Build the mobile app with Cloud SQL as the backend
- B. Build the mobile app with Cloud Storage as the backend
- C. Build the mobile application with Firebase as the backend
- D. Build the mobile app with Cloud Spanner as the backend

Answer: C

Explanation:

Firebase/Firestore is easy to build and is suitable for user information that could vary in nature.

NEW QUESTION 230

- (Topic 2)

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

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

Answer: B

Explanation:

You can send notification messages using the Notifications composer in the Firebase console. Though this does not provide the same flexibility or scalability as sending messages with the Admin SDK or the HTTP and XMPP protocols, it can be very useful for testing or for highly targeted marketing and user engagement. The Firebase console provides analytics-based A/B testing to help refine and improve marketing messages.

After you have developed logic in your app to receive messages, you can allow non-technical users to send messages per the instructions on the Notifications page in the Firebase Help Center.

NEW QUESTION 232

- (Topic 2)

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

- A. Circular Scaling takes a round-robin approach to allocate and destroy VMs.
- B. Triangular Scaling takes an automated average of Cost, Effort, and Time.
- C. Vertical Scaling
- D. Horizontal Scaling

Answer: D

Explanation:

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

NEW QUESTION 237

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 238

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 241

- (Topic 2)

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

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

Answer: BCE

Explanation:

Because GCP provides three options for paid support which are Standard, Enhanced and Premium. Basic Support is included with your Google Cloud subscription which cover only Case, phone, and chat support for billing issues only Reference link- <https://cloud.google.com/support>

NEW QUESTION 245

- (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 249

- (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 251

- (Topic 2)

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

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

Answer: D

Explanation:

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

-> Google Cloud Armor security policies are available only for backend services behind an external HTTP(S) load balancer. The load balancer can be in Premium Tier or Standard Tier.

-> Google Cloud Armor security policies and IP DENY lists and ALLOW lists are available only for HTTP(S) load balancing.

Reference link- <https://cloud.google.com/armor/docs/security-policy-overview>

NEW QUESTION 254

- (Topic 2)

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

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

Answer: C

Explanation:

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

Lift and shift

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

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

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

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

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

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

NEW QUESTION 258

- (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 260

- (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-slals-and-slos>

NEW QUESTION 261

- (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 264

- (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 268

- (Topic 2)

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

Multi-regions

Multi-Region Name	Multi-Region Description
ASIA	Data centers in Asia
EU	Data centers within member states of the European Union*
US	Data centers in the United States

- A. Cloud Storage with a single region that is known to be within the European U
- B. Cloud Filestore is connected to virtual machines which are guaranteed to be within the European U
- C. Cloud Storage with the multi-region option of European U
- D. Cloud Storage with the dual-region option of European U

Answer: C

Explanation:

Multi-region option will use multiple datacenters that are within the European Union. More regions will also help with lower latency since users are spread across the European U.

<https://cloud.google.com/storage/docs/locations#considerations>

NEW QUESTION 269

- (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 automatically 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 273

- (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 275

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 276

- (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 278

- (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 281

- (Topic 2)

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

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

Answer: D

Explanation:

These are the feature of Infrastructure as a Service (IaaS) It's economical

Because IaaS resources are used on demand and enterprises only have to pay for the compute, storage, and networking resources that are actually used, IaaS costs are fairly predictable and can be easily contained and budgeted for.

It's efficient

IaaS resources are regularly available to businesses when they need them. As a result, enterprises reduce delays when expanding infrastructure and, alternatively, don't waste resources by overbuilding capacity.

It boosts productivity

Because the cloud provider is responsible for setting up and maintaining the underlying physical infrastructure, enterprise IT departments save time and money and can redirect resources to more strategic activities.

It's reliable

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

It's scalable

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

It drives faster time to market

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

NEW QUESTION 285

- (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 286

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 289

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 292

- (Topic 2)

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

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

Answer: D

Explanation:

Cloud Spanner:

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

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

NEW QUESTION 296

- (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 299

- (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:

Firestore Cloud Messaging:

Firestore 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 Firestore 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 302

- (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 304

- (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 309

- (Topic 2)

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

ment. What would be your recommendation?

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

Answer: B

Explanation:

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

Instance groups

[Send feedback](#)

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

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

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

[Learn more](#)

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

NEW QUESTION 311

- (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 314

- (Topic 2)

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

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

Answer: A

Explanation:

Containers and Kubernetes are ideal for the kind of requirement mentioned here - separate microservices that need to scale independently.

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

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

NEW QUESTION 315

- (Topic 2)

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

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

Answer: BC

Explanation:

Customer-managed encryption keys for Cloud BigTable.

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

If you have specific compliance or regulatory requirements related to the keys that protect your data, you can use customer-managed encryption keys (CMEK) for BigTable. Instead of Google managing the encryption keys that protect your data, your BigTable instance is protected using a key that you control and manage in Cloud Key Management Service (Cloud KMS).

Features

Security: CMEK provides the same level of security as Google's default encryption but provides more administrative control.

Data access control: Administrators can rotate, manage access to, and disable or destroy the key used to protect data at rest in BigTable .

Auditability: All actions on your CMEK keys are logged and viewable in Cloud Logging. Comparable performance: BigTable CMEK-protected instances offer comparable performance to BigTable instances that use Google default encryption.

Flexibility: You can use the same CMEK key in multiple projects or instances or you can use separate keys, depending on your business needs.

NEW QUESTION 317

- (Topic 2)

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

- A. Use Cloud Vision API - write code to identify the text blocks, copy the data, and store it
- B. Use TensorFlow - write code that will identify the type of visa and the bounding text block
- C. Copy the data and then store it.
- D. Use AutoML - upload other images of visas and run the model creation process which will automatically identify the visas
- E. Use Data Labeling service - outsource the work of marking and extracting the information to others.

Answer: A

Explanation:

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

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

NEW QUESTION 322

- (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 323

- (Topic 2)

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

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

Answer: C

Explanation:

Lending DocAI is a pre-packaged AI solution that speeds "up the mortgage workflow processes to easily process loans and automate document data capture, while ensuring the accuracy and breadth of different documents (e.g., tax statements and asset documents)."

<https://cloud.google.com/solutions/lending-doc-ai>

NEW QUESTION 324

- (Topic 2)

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

- A. A type of machine learning model for distinguishing among two or more discrete value
- B. E.
- C. "book", "car".
- D. A type of machine learning model is a meta-model maker, which classifies algorithms based on the quality of their output.
- E. A type of machine learning model that outputs continuous (typically, floating-point) value
- F. E.
- G. the predicted price of the house is \$120,000.
- H. A type of classic model approach that is less used today and which has been replaced by the regression model.

Answer: A

Explanation:

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

NEW QUESTION 327

- (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 331

- (Topic 2)

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

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

Answer: B

Explanation:

SaaS – Software as a Service (SaaS) provides you a complete product that is run and managed by the service provider. You worry only about using the software and not about infrastructure.

SaaS provides the lowest level of flexibility and management control over the infrastructure. (Example: Google Gsuite and MS O365)

NEW QUESTION 334

- (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 account 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 339

- (Topic 2)

You have a well established development and operations team. Your teams were managing the entire software delivery/deployment cycle on-premise. When migrating to the cloud, you want to continue having this approach. Which is the ideal option for you?

- A. PaaS - Platform as a Service
- B. SaaS - Software as a Service

- C. IDaaS - Identity as a Service
- D. IaaS - Infrastructure as a Service

Answer: D

Explanation:

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

NEW QUESTION 342

- (Topic 2)

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

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

Answer: B

Explanation:

Cloud Spanner:

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

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

Automatic sharding

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

Strong transactional consistency

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

Regional and multi-regional configurations

No matter where your users may be, apps backed by Cloud Spanner can read and write up-to-date strongly consistent data globally. Additionally, when running a multi-region instance, your database is able to survive a regional failure, and offers industry-leading 99.999% availability.

Online schema changes with no downtime

Cloud Spanner users can make a schema change, whether it's adding a column or adding an index while serving traffic with zero downtime. Hence you now have the flexibility to adapt your database to your business needs without compromising on the availability of your application.

NEW QUESTION 343

- (Topic 2)

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 344

- (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 345

- (Topic 1)

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

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

Answer: A

Explanation:

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

Reference link:-

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

NEW QUESTION 346

- (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 348

- (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 353

- (Topic 1)

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

Where should your organization locate this virtual machines?

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

Answer: B

Explanation:

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

NEW QUESTION 355

- (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 356

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 360

- (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 363

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 365

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 366

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 367

- (Topic 1)

What conditions be true if a VM interface wants to send packets to the external IP addresses of Google APIs and services using Private Google Access?

- A. VM interface does not have an external IP address assigned.
- B. VM interface is connected to a subnet where Private Google Access is disabled
- C. Both A and B
- D. None of the Above.

Answer: A

Explanation:

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

- The VM interface is connected to a subnet where Private Google Access is enabled.
 - The VPC network that contains the subnet meets the network requirements for Google APIs and services.
 - The VM interface does not have an external IP address assigned.
 - The source IP address of packets sent from the VM matches the VM interface's primary internal IP address or an internal IP address from an alias IP range.
- A VM with an external IP address assigned to its network interface doesn't need Private Google Access to connect to Google APIs and services. However, the VPC network must meet the requirements for accessing Google APIs and services.

NEW QUESTION 371

- (Topic 1)

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

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

Answer: B

Explanation:

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

Introduction to the Organization Policy Service

Send feedback

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

Benefits

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

References link:

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

NEW QUESTION 376

- (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?

- Use a committed use discount to reserve a very powerful virtual machine
- Purchase one very powerful physical computer
- Start a very powerful virtual machine without using a committed use discount
- 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 379

- (Topic 1)

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

- Elasticity
- Availability
- Broad Network Access
- Multi-Tenancy

Answer: D

Explanation:

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

NEW QUESTION 380

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 382

- (Topic 1)

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

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

Answer: D

Explanation:

Anthos

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

Try it free

Contact sales

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

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

NEW QUESTION 387

- (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 389

- (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 392

- (Topic 1)

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

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

Answer: CD

Explanation:

General Availability is the stage where SLAs apply.

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

NEW QUESTION 395

- (Topic 1)

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

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

Answer: D

Explanation:

Option A is true

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

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

Responsibilities Data, including:

- Security and encryption
- Backups

Software and applications, including:

- Installation
- Configuration
- Upgrades and patching

Operating system and any hypervisor, including:

- Configuration changes
- Upgrades and patching Server clusters, including:
 - Installation
 - Configuration
 - Maintenance Licensing

Option B & C is also true.

With Bare Metal Solution, Google Cloud provides and manages the core infrastructure, the net-work, the physical and network security, and hardware monitoring capabilities in an

environment from which you can access all of the Google Cloud services. The core infrastructure includes secure, controlled-environment facilities, and power.

The Bare Metal Solution also includes the provisioning and maintenance of the custom, sole-tenancy hardware with local SAN, and smart hands support.

The network, which is managed by Google Cloud includes a low-latency Cloud Interconnect con-nection into the customer Bare Metal Solution environment.

The available Google Cloud services include private API access, management tools, support, and billing.

NEW QUESTION 397

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your Cloud-Digital-Leader Exam with Our Prep Materials Via below:

<https://www.certleader.com/Cloud-Digital-Leader-dumps.html>