



Google

Exam Questions Cloud-Digital-Leader

Google Cloud Digital Leader exam

About Exambible

Your Partner of IT Exam

Found in 1998

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

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

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

NEW QUESTION 1

- (Topic 1)

Your organization is defining the resource hierarchy for its new application in Google Cloud. You need separate development and production environments. The production environment will be deployed in Compute Engine in two regions. Which structure should your organization choose?

- A. Create a single project for all environment
- B. Use labels to segregate resources by environment.
- C. Create a single project for all environment
- D. Use tags to segregate resources by environment.
- E. Create one project for the development environment and one project for the production environment.
- F. Create two projects for the development environment and two projects for the production environment (one for each region).

Answer: C

Explanation:

Many organizations have separate development and production environments so they can build and test new features without disturbing production traffic. In Optimizely, you can create separate projects for each environment to help with governance.

With separate development and production projects, your organization can safely build and QA experiments and Personalization campaigns in a development environment before deploying to production. This approach allows multiple stakeholders in your organization to act as gatekeepers for running new experiments in production.

Set up projects

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

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

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

NEW QUESTION 2

- (Topic 1)

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

- A. Keep the credentials separate for each application to reduce the blast radius in case of any issues.
- B. Use an external identity provider that is famous and popular like Facebook or Twitter; that way, even your vendors and contractors will have an account there.
- C. Allow all users, especially contractors and vendors, to bring their own identities, like those at gmail.com.
- D. Use an IDaaS (Identity as a Service) product that can centrally manage authentication and authorization for the applications.

Answer: D

Explanation:

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

NEW QUESTION 3

- (Topic 1)

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 4

- (Topic 1)

Your organization meant to purchase a 3-year Committed Use Discount, but accidentally purchased a 1-year Committed Use Discount instead. What should your organization do?

- A. Contact your financial institution.

- B. Contact Trust and Safety.C Contact Cloud Billing Support.
C. Contact Technical Support.

Answer: C

Explanation:

★ **Important:** Once you enable discount sharing through the console, you can only disable it with the assistance of **Cloud Billing support**. If you disable discount sharing with the assistance of Cloud Billing support, all committed use discounts revert to the default setting of applying only to the projects through which with they were purchased. The reverted setting becomes effective at the beginning of the following month.

Combining reservations with commitments

A committed use discount provides a 1- or 3-year discounted price agreement, but it does not reserve capacity in a specific zone. A reservation ensures that capacity is held in a specific zone even if the reserved VMs are not running. By combining a reservation with a commitment, you get discounted, reserved resources.

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

NEW QUESTION 5

- (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 enhancedsupport ensuring the long term success of your SecureAuth SaaS IAM deployment

NEW QUESTION 6

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 7

- (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 8

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 9

- (Topic 1)

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

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

Answer: A

Explanation:

Cloud Memorystore is an in-memory database that has sub-millisecond latency. This is ideal for caching application data that also changes once in a while.

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

NEW QUESTION 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 10

- (Topic 1)

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

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

Answer: B

Explanation:

Migrate for Compute Engine enables you to lift and shift workloads at scale to Google Cloud Compute Engine with minimal changes and risk.

Reference: <https://dataintegration.info/simplify-vm-migrations-with-migrate-for-compute-engine-as-a-service>

NEW QUESTION 13

- (Topic 1)

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

How should you meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 14

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 18

- (Topic 1)

Your organization is developing a plan for migrating to Google Cloud.

What is a best practice when initially configuring your Google Cloud environment?

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

Answer: B

Explanation:

The Organization resource is the root node of the Google Cloud resource hierarchy and all resources that belong to an organization are grouped under the organization node. This provides central visibility and control over every resource that belongs to an organization.

Reference link- <https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy>

NEW QUESTION 23

- (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 27

- (Topic 1)

A company with its own private data center has called you in for help with their disaster recovery planning. News of multiple ransomware attacks has made them very anxious. They want to make they are well prepared for such an eventuality. Which of these would be good recommendations?

- A. It is better to have redundancy; so, set up another private data center nearby so that you can quickly go over in case of an emergency.
- B. It is better to have redundancy; use one or many of the Google Cloud datacenters as a backup location.
- C. The one data center is enough, as long as the data is encrypted; attackers won't be able to read the data.
- D. The one data center is enough as long as you regularly back up data and save it in another place in the same DC.

Answer: B

Explanation:

A single data center is vulnerable. So any option involving that is not good. Reference Link:- <https://www.coresite.com/blog/data-center-redundancy>

NEW QUESTION 32

- (Topic 1)

Your customer is making a decision on whether to move to Google Cloud. Their key concern is about 10,000 VMs that are part of their IT infrastructure used across more than 110 applications. They are apprehensive of too many changes at this stage. They want to get to Google Cloud in the easiest way possible with minimal disruption. What option would you recommend for them?

- A. Use Migrate for Anthos
- B. Lift and shift the VMs to serverless options like App Engine Flex.
- C. Re-architect on-prem to use Kubernetes and then slowly extend and bridge the on-prem data center to the Google Cloud data center.
- D. Use Migrate for Compute

Answer: D

Explanation:

Migrate for Compute Engine's advanced replication migration technology copies instance data to Google Cloud in the background with no interruptions to the source workload that's running.

Cloud migration creates a lot of questions. Migrate for Compute Engine by Google Cloud has the answers. Whether you're looking to migrate one application from on-premises or one thousand enterprise-grade applications across multiple data centers, Migrate for Compute Engine gives any IT team, large or small, the power to migrate their workloads to Google Cloud.

Watch the video to your right to hear what one of our customers, Rackspace Technology, thinks about Migrate for Compute Engine's speed and ease of use.

<https://cloud.google.com/migrate/compute-engine>

NEW QUESTION 34

- (Topic 1)

Your organization wants to predict the behavior of visitors to its public website. To do that, you have decided to build a machine learning model. Your team has database-related skills but only basic machine learning skills, and would like to use those database skills.

Which Google Cloud product or feature should your organization choose?

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

Answer: A

Explanation:

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

NEW QUESTION 37

- (Topic 1)

Your organization wants to optimize its use of Google Cloud's discounts on virtual machine-based workloads. You plan to use 200 CPUs constantly for the next 3 years, and you forecast that spikes of up to 300 CPUs will occur approximately 30% of the time. What should you choose?

- A. 1-year committed use discount for 200 CPUs
- B. 3-year committed use discount for 300 CPUs
- C. 3-year committed use discount for 200 CPUs
- D. Regular pay-as-you-go pricing

Answer: C

Explanation:

you can get a 57% discount by agreeing to commit to a 3-year contract. Any usage over the commitment will just be billed at the standard rate. Since they only need 300 CPUs 30% of the time, will pick answer C so that we are not paying usage off 300 CPUs all of the time. This gives us a discount of 57% for 200 CPU's, huge cost savings.

NEW QUESTION 38

- (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 39

- (Topic 1)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 43

- (Topic 1)

Your company has been using a shared facility for data storage and will be migrating to Google Cloud. One of the internal applications uses Linux custom images that need to be migrated.

Which Google Cloud product should you use to maintain the custom images?

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

Answer: B

Explanation:

Reference: <https://cloud.google.com/compute/docs/images/create-delete-deprecate-private-images>

A custom image is a boot disk image that you own and control access to. Use custom images for the following tasks:

Import a virtual disk to Compute Engine from your on-premises environment or from VMs that are running on your local workstation or on another cloud platform.

You can manually import boot disk images to Compute Engine, but one disk at a time.

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

<https://cloud.google.com/compute/docs/images>

NEW QUESTION 46

- (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 49

- (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 50

- (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 54

- (Topic 1)

Your organization is developing a mobile app and wants to select a fully featured cloud- based compute platform for it. Which Google Cloud product or feature should your organization use?

- A. Google Kubernetes Engine
- B. Firebase
- C. Cloud Functions
- D. App Engine

Answer: B

Explanation:

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

Firebase is Google's mobile development platform that empowers you to quickly build and grow your app

NEW QUESTION 57

- (Topic 1)

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

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

Answer: A

Explanation:

BigQuery ML lets you create and execute machine learning models in BigQuery using standard SQL queries.

Reference: <https://cloud.google.com/bigquery-ml/docs/introduction#:~:text=BigQuery%20ML%20lets%20you%20create,the%20need%20to%20move%20data>

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

<https://cloud.google.com/bigquery-ml/docs/introduction>

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

Answer: D

Explanation:

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

NEW QUESTION 63

- (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 68

- (Topic 3)

An organization needs frequent access to only a subset of their data. They want to reduce costs by depositing the rest of their data across Nearline Coldlme and Archive repositories

Which Google Cloud product should the organization use?

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

Answer: D

Explanation:

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

NEW QUESTION 73

- (Topic 3)

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

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

Answer: D

NEW QUESTION 77

- (Topic 3)

An organization wants to collect metrics and metadata from their cloud applications and put them into dashboards. Which Google Cloud tool should they use?

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

Answer: A

Explanation:

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

NEW QUESTION 81

- (Topic 3)

An organization needs to search an application's source code to identify a potential issue. The application is distributed across multiple containers. Which Google Cloud product should the organization use?

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

Answer: B

Explanation:

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

NEW QUESTION 85

- (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 87

- (Topic 3)

How can a streaming service meet global compliance requirements using the cloud?

- A. By automatically encrypting personally identifiable information
- B. By obtaining a business license to operate in a new market
- C. By allowing users to disable two-factor authentication
- D. By adhering only to data policies of the country in which the head office is registered

Answer: A

NEW QUESTION 92

- (Topic 3)

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

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

Answer: B

NEW QUESTION 94

- (Topic 3)

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

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

Answer: A

NEW QUESTION 96

- (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 98

- (Topic 3)

An organization provides a loyalty program for its customers. It recently partnered with other businesses so that customers can get loyalty points at a range of other stores.

Why should the organization use application programming interfaces (APIs)?

- A. To migrate all partner data for disaster recovery
- B. To analyze and publish loyalty program statistics to a dashboard
- C. To personalize recommendations for loyalty card users
- D. To connect third-party systems to ensure up-to-date information

Answer: D

NEW QUESTION 99

- (Topic 3)

How does Google Cloud ensure that customer data remains secure and private when at rest?

- A. By aggregating training data for customers within each industry
- B. By automatically locking files containing suspicious code
- C. By auditing platform privacy practices against industry standards
- D. By providing privacy reviews for critical customer applications

Answer: C

Explanation:

Google Cloud commitment to keep the data secure and private:

- * 1. Org owns the data and not Google
- * 2. Google does not sell data to 3rd parties
- * 3. All customer data is encrypted by default
- * 4. Google Cloud guards insider against your data
- * 5. No backdoor access to any govt. entity
- * 6. Google's privacy practices are audited against international standards

NEW QUESTION 103

- (Topic 3)

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

- A. The inaccessibility of their data due to perimeter security
- B. The overreliance on platform as a service
- C. The time it takes their serverless compute function to scale

D. The cost of provisioning hardware for peak usage

Answer: D

Explanation:

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

NEW QUESTION 107

- (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 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)

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

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

Answer: D

Explanation:

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

NEW QUESTION 117

- (Topic 3)

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

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

Answer: B

NEW QUESTION 120

- (Topic 3)

How does switching from on-premises to the cloud help organizations gain value over time?

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

Answer: A

NEW QUESTION 123

- (Topic 3)

An organization wants to build an entirely new infrastructure and applications in the cloud. Which application modernization approach should the organization use?

- A. Move the application to the cloud, and then change it.
- B. Change their application, and then move it to the cloud.

- C. Invent in greenfield.
- D. Invent in brownfield.

Answer: C

Explanation:

A Greenfield approach is a brand-new implementation , where companies then add their needed configurations and customizations. This approach provides a clean slate to start from, does not carry over needless customizations and technical debt, and provides a solid foundation for business process re-engineering. A greenfield deployment is the design, installation and configuration of computer infrastructure where none existed before, for example, in a new office. In contrast, a brownfield deployment is an upgrade or addition to existing infrastructure using legacy components.

NEW QUESTION 127

- (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 131

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 133

- (Topic 3)

An organization needs a platform to create custom end-to-end artificial intelligence models. Which Google Cloud product or service should the organization use?

- A. Dataproc
- B. Compute Engine
- C. Recommendations AI
- D. Vertex AI

Answer: D

Explanation:

Recommendations AI enables you to build an end-to-end personalized recommendation system based on state-of-the-art deep learning ML models, without a need for expertise in ML or recommendation systems. With Vertex AI, both AutoML training and custom training are available options. Whichever option you choose for training, you can save models, deploy models, and request predictions with Vertex AI. <https://cloud.google.com/vertex-ai>

NEW QUESTION 138

- (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 142

- (Topic 3)

Why do organizations often struggle to scale their on-premises application infrastructure?

- A. Scaling compute instances could breach compliance and/or regulation
- B. Increasing compute capacity is time-consuming and costly

- C. Their serverless compute functions struggle to meet the demand
- D. Their multi-cloud architecture is complex and expensive

Answer: B

NEW QUESTION 145

- (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 149

- (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 151

- (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 156

- (Topic 3)

An organization wants to transform multiple types of structured and unstructured data in the cloud from various sources. The data must be readily accessible for analysis and insights.

Which cloud data storage system should the organization use?

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

Answer: D

Explanation:

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

NEW QUESTION 157

- (Topic 3)

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

What should the organization use?

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

Answer: B

NEW QUESTION 159

- (Topic 3)

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

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

Answer: A

NEW QUESTION 161

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 162

- (Topic 3)

An organization wants to upskill their IT staff. How can they do this in a transformational way?

- A. Prioritize training current employees instead of hiring new recruits with cloud experience.
- B. Prioritize giving privileged access to third-party partners and contractors to fill IT knowledge gaps.
- C. Create a culture of self-motivated, isolated learning with official training materials.
- D. Create a culture of continuous peer-to-peer learning with official training materials.

Answer: D

NEW QUESTION 167

- (Topic 3)

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

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

Answer: B

NEW QUESTION 170

- (Topic 3)

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 174

- (Topic 3)

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

- A. It analyzes unstructured data which can then be accessed in multiple regions

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

Answer: B

Explanation:

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

NEW QUESTION 176

- (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 177

- (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 181

- (Topic 3)

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

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

Answer: B

NEW QUESTION 182

- (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 meets their service level objective (SLO) of 99 999% ("five nines") How much downtime do their end users experience per year?

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

Answer: A

NEW QUESTION 191

- (Topic 3)

What does Cloud Debugger help an organization do?

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

Answer: B

Explanation:

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

NEW QUESTION 192

- (Topic 2)

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

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

Answer: A

Explanation:

Because Dataproc is used to run Hadoop/Spark workloads

NEW QUESTION 193

- (Topic 2)

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

- 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 innova-tion.
- D. Third-party systems are less capable of addressing an organization's security re- quirements.

Answer: B

Explanation:

Because unsecured third-party systems are a cybersecurity threat.

NEW QUESTION 195

- (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 200

- (Topic 2)

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

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

Answer: BC

NEW QUESTION 205

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 207

- (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 211

- (Topic 2)

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

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

Answer: AB

NEW QUESTION 215

- (Topic 2)

A large organization is struggling to manage their cloud costs effectively. They want to increase visibility into cloud costs. Which cost management approach should the organization use?

- A. Establish a partnership between finance, technology, and business teams.
- B. Appoint a single person to monitor cloud spending across the organization.
- C. Review any cloud spending that exceeds the organization's error budget.
- D. Increase monitoring of on-premises infrastructure and services.

Answer: A

Explanation:

Because cross-team partnerships are part of the visibility cost management strategy.
https://wa.aws.amazon.com/wat/question.COST_1.en.html

NEW QUESTION 216

- (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 219

- (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 222

- (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 covers only Case, phone, and chat support for billing issues only

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

NEW QUESTION 226

- (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:

Firebase Crashlytics:

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

Firebase 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 228

- (Topic 2)

Your company has signed up with a cloud provider and you will be using storage and virtual machines with the provider. The provider has provided your organization some expectations for what the service should perform at. What type of agreement provides a guarantee of a certain level of service such as “Uptime”?

- A. Performance Agreement
- B. Interconnection Agreement
- C. Warranty
- D. Service Level Agreement

Answer: D

Explanation:

Service Level Agreement (SLA)

A service level agreement (SLA) is a contract between a service provider (either internal or external) and the end user that defines the level of service expected from the service provider. Some common SLA’s are uptime, Response Time, etc.

NEW QUESTION 230

- (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	
Multicloud 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 231

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 232

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 236

- (Topic 2)

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

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

Answer: B

Explanation:

Coldline storage is a very low cost highly durable storage service for data archiving, online backup, and disaster recovery. Coldline storage is the best choice for data

that you plan to access at most once a year due to its slightly lower availability, 90 day minimum storage duration cost for data access, and higher per operation costs. Nearline and Coldline are for backup and archival storage and having the highest availability for both with 99.9 percent.

NEW QUESTION 237

- (Topic 2)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 238

- (Topic 2)

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

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

Answer: D

Explanation:

Flavors of Apigee

Apigee comes in the following flavors:

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

services.

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

NEW QUESTION 241

- (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 242

- (Topic 2)

A customer in the European Union region is very clear that their data should not go outside the Eu-ropean 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 244

- (Topic 2)

Which of the following is/are core storage options available on the Google Cloud Platform?

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

Answer: D

Explanation:

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

NEW QUESTION 245

- (Topic 2)

The government has mandated that companies in a particular section of healthcare must retain all the data they collect for a period of 10 years in case an audit needs to be done. Your client, who is in that industry, needs to follow regulations. In addition, your client wants to do an analysis of the data quite frequently in the first year. They also don't want to be liable for any data beyond year 10. What would recommend for your customer?

- A. Use Cloud Storage with nearline storage in year one and Coldline storage thereaf-te
- B. Use Object lifecycle management to move between storage types and delete them after 10 years.
- C. Use Cloud Storage with standard storage in year one and Coldline storage there-afte

- D. Set a Cloud Scheduler trigger for 1 year to change storage types and 10 years to delete the data.
- E. Use Cloud Storage with standard storage in year one and archival storage thereafter.
- F. Use Object lifecycle management to move between storage types and delete them after 10 years.
- G. Use Cloud Storage with standard storage in year one and Coldline storage thereafter.
- H. Set a Cloud Tasks to trigger for 1 year to change storage types and 10 years to delete the data.

Answer: C

Explanation:

Cloud storage supports Object Lifecycle Management. To support common use cases like setting a Time to Live (TTL) for objects, retaining noncurrent versions of objects, or "downgrading" storage classes of objects to help manage costs, Cloud Storage offers the Object Lifecycle Management feature. Standard storage is recommended for frequently accessed data and Archive for data accessed less than once a year. Nearline, Coldline, and Archive offer ultra-low-cost, highly-durable, highly available archival storage. For data accessed less than once a year, Archive is a cost-effective storage option for the long-term preservation of data. Coldline is also ideal for cold storage—data your business expects to touch less than once a quarter. For warmer storage, choose Nearline: data you expect to access less than once a month, but possibly multiple times throughout the year.

NEW QUESTION 247

- (Topic 2)

Firebase Hosting provides the following services-

- A. Dynamic content
- B. Static content.
- C. Microservices.
- D. All of the Above.

Answer: D

Explanation:

Firebase Hosting- Firebase Hosting provides fast and secure hosting for your web app, static and dynamic content, and microservices. Firebase Hosting is production-grade web content hosting for developers. With a single command, you can quickly deploy web apps and serve both static and dynamic content to a global CDN (content delivery network). You can also pair Firebase Hosting with Cloud Functions or Cloud Run to build and host microservices on Firebase. Key capabilities of Firebase Hosting: Serve content over a secure connection:- The modern web is secure. Zero-configuration SSL is built into Firebase Hosting, so content is always delivered securely. Host static and dynamic content plus microservices:- Firebase Hosting supports all kinds of content for hosting, from your CSS and HTML files to your Express.js microservices or APIs. Deliver content fast: Each file that you upload is cached on SSDs at CDN edges around the world and served as gzip or Brotli. We auto-select the best compression method for your content. No matter where your users are, the content is delivered fast.

NEW QUESTION 251

- (Topic 2)

You are a program manager in a company you need to submit a bare metal solution order for a secure, high performance connection with a low-latency network fabric. What network information you need to submit the order to Bare Metal Solutions.

- A. IP Ranges for example Client IP Address range used for communication between your Google Cloud and Bare Metal Solution environments.
- B. Google Cloud Project Id that you are using with your bare metal solution environment.
- C. Total number of VLANs you need in your Bare Metal Solution Environment.
- D. All of the above

Answer: D

Explanation:

What Bare Metal Solution provides Bare Metal Solution is a managed solution that provides purpose-built HPE or Atos bare-metal servers in regional extensions that are connected to Google Cloud by a managed, high-performance connection with a low-latency network fabric. With Bare Metal Solution, Google Cloud provides and manages the core infrastructure, the network, 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 custom, sole-tenancy servers with local SAN, and smart hands support. The network, which is managed by Google Cloud, includes a low-latency Partner Interconnect connection into the customer Bare Metal Solution environment. The available Google Cloud services include private API access, management tools, support, and billing.

NEW QUESTION 256

- (Topic 2)

What type of cloud computing service provides raw compute, storage, and network, organized in ways that are familiar to physical data centers?

- A. Database as a Service.
- B. Platform as a Service.
- C. Infrastructure as a Service.
- D. Software as a Service.

Answer: C

Explanation:

What is Infrastructure as a service : IaaS (infrastructure as a service) is a computing model that offers resources on-demand to businesses and individuals via the cloud. IaaS is attractive because acquiring computing resources to run applications or store data the traditional way requires time and capital. Enterprises must purchase equipment through procurement processes that can take months. They must invest in physical spaces: typically specialized rooms with power and cooling. And

after deploying the systems, enterprises need IT, professionals, to manage them.

All this is challenging to scale when demand spikes or the business grows. Enterprises risk running out of capacity or overbuilding and ending up with infrastructure that suffers from low utilization.

These challenges are why IaaS use is steadily growing. Learn more about Compute Engine, Cloud Storage, etc.

NEW QUESTION 261

- (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 265

- (Topic 2)

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

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

Answer: B

Explanation:

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

<https://cloud.google.com/solutions/web-app-and-api-protection>

NEW QUESTION 266

- (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 them
- 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 features 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 270

- (Topic 2)

You are consulting for a client who is migrating to Google Cloud. They presently have a matrix organization. Their IT environments were managed around projects. Each team had multiple projects. All the projects had a flat structure under the company. What would you advise them when planning for the move?

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

Answer: A

Explanation:

Folders for a related group of projects are the recommended approach.

-> A flat structure under the organization node is possible on Google Cloud, but it is not recommended. It becomes tougher to manage.

-> Projects cannot have sub-projects; there can only be resources within Projects.

Reference link- <https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy>

NEW QUESTION 273

- (Topic 2)

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

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

Answer: AD

Explanation:

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

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

Reference link- https://en.wikipedia.org/wiki/High_availability

NEW QUESTION 276

- (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 280

- (Topic 2)

In discussions with a prospective customer who wants to move to Google Cloud to make use of the latest, scalable technologies available therein, you learn that there are very strict regulations concerning the storage of data. They only have the approval to store it in their current private data center. What would you advise them?

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

Answer: A

Explanation:

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

NEW QUESTION 285

- (Topic 2)

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

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

Answer: CDE

Explanation:

Sustained, Committed and Preemptible

vCPU and memory usage for each of these categories can receive discounts VM vCPU and memory usage for each of these categories can receive discounts

Sustained-use discounts—Google offers up to 30% off for workloads that run for most of the billing month on GCP services.

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

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

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

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

NEW QUESTION 290

- (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 291

- (Topic 2)

Your company has a requirement to run manual tests on their web products for UX research before it is released to end customers. The people who will do the tests are external to the company. They will either use their own Gmail id or be given temporary email ids using the applications and recording their inputs in another app. The UX testing is done in the last week of the month. Each month the UX testers could be different. How should the IT team manage the users?

- A. Since the app is anyways going to be public, create permanent credentials for the UX testers that they can conveniently use each time.
- B. It would be a security issue to have users come and go
- C. Recommend that the testers be permanently hired to plug the vulnerability issue.
- D. It would be a security issue to have users come and go
- E. Recommend that the testers be permanently hired to plug the vulnerability issue.
- F. Create a Group with the permissions required to do the test and record their input
- G. When users arrive each week, add them to the group and after the testing period, remove them from the group.

Answer: D

Explanation:

Groups are convenient to use for this requirement. Permissions to the group are automatically inherited by the members of the group. Adding and removing UX testers from the group will grant and remove permissions.

NEW QUESTION 293

- (Topic 2)

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

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

Answer: B

Explanation:

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

Choose a Google Cloud bare metal migration strategy

[Send feedback](#)

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

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

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

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

NEW QUESTION 296

- (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 300

- (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 303

- (Topic 2)

Cloud Data Loss Prevention (DLP) is a fully managed service designed to help discover, classify, and protect the most sensitive data. DLP provides three key features (Select Three Answers)

- A. Classification
- B. De-identification
- C. De-classification

- D. Inspection
- E. Reinspection

Answer: ABD

Explanation:

Classification. De-classification and Inspection

Classification is the process to inspect the data and know what data we have, how sensitive it is, and the likelihood. Inspection and classification happen here.

De-identification is the process of removing, masking, replacing information from data.

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

NEW QUESTION 306

- (Topic 2)

With respect to the Core Feature of Standby Instances of Cloud SQL which one of the options is correct.?

- A. The standby instance is used in high availability to replace the primary instance when failover occur
- B. The standby instance appears in the Google Cloud Console but does not get billed
- C. When failover occurs, connections to the primary instance need to be manually transferred to the standby instance.
- D. The standby instance is used in high availability to replace the primary instance when failover occur
- E. The standby instance appears in the Google Cloud Console but does not get billed
- F. When failover occurs, connections to the primary instance are automatically transferred to the standby instance.
- G. The standby instance is used in high availability to replace the primary instance when failover occur
- H. The standby instance doesn't appear in the Google Cloud Console
- I. When failover occurs, connections to the primary instance are automatically transferred to the standby instance.
- J. None of the Above.

Answer: C

Explanation:

The standby instance is used in high availability to replace the primary instance when failover occurs. The standby instance doesn't appear in the Google Cloud Console. When failover occurs, connections to the primary instance are automatically transferred to the standby instance.

Cloud SQL Key Terms: Cloud SQL instance

A Cloud SQL instance corresponds to one virtual machine (VM). The VM includes the database instance and accompanying software containers to keep the database instance up and running.

Database instance

A database instance is the set of software and files that operate the databases: MySQL, PostgreSQL or SQL Server.

High availability

Cloud SQL instances using high availability (HA) provide greater reliability than non-HA instances.

HA in Cloud SQL works by having two synchronized instances: a primary instance and a standby instance. Each instance has exactly one VM. Each instance is in a different zone in the same region.

Failover

A failover is when Cloud SQL switches serving from the original primary instance to the standby instance.

Autofailover is a mechanism that automatically triggers failover when a Cloud SQL instance didn't issue a heartbeat in the previous interval.

Standby instances

The standby instance is used in high availability to replace the primary instance when failover occurs. The standby instance doesn't appear in the Google Cloud Console. When failover occurs, connections to the primary instance are automatically transferred to the standby instance.

Clone

When you clone a Cloud SQL instance, you create a new instance that is a copy of the source instance, but is completely independent. After cloning is complete, changes to the

source instance are not reflected in the clone, and changes in the clone are not reflected in the source instance.

Replication

Replication is the ability to create copies of a Cloud SQL instance or an on-premises database, and offload work to the copies. The main reason for using replication is to scale the use of data in a database without degrading performance on the primary instance. Read replica

The read replica is an exact copy of the primary instance. Data and other changes on the primary instance are updated in almost real time on the read replica.

Send your write transactions to the primary instance, and your read requests to the read replica. The read replica processes queries, read requests, and analytics traffic, thus reducing the load on the primary instance.

Source server

Replication copies transactions from a primary instance to one or more read replicas. The primary instance is also called the source server. The source server can be a Cloud SQL primary instance, or a server outside of Google Cloud, such as an on-premises server or a server running in a different cloud. If the source server is outside of Google Cloud, we call it Replication from an external server.

Cloud SQL Auth proxy client

The Cloud SQL Auth proxy client is open source software maintained by Cloud SQL. It connects to a companion process, the Cloud SQL Auth proxy server, running on your Cloud SQL instance. You run the Cloud SQL Auth proxy client on your own servers. The Cloud SQL Auth proxy client can be used to establish a secure SSL/TLS connection to the database instance, and/or to avoid having to open the firewall. Authentication is done through Identity and Access Management (IAM).

NEW QUESTION 307

- (Topic 2)

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

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

Answer: B

Explanation:

Cloud BigTable

Key features

High throughput at low latency

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

Cluster resizing without downtime

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

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

NEW QUESTION 312

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 314

- (Topic 2)

A bank wants to track the success of their existing ATM network, which has been modernized with APIs to instantly notify customers about their transfers. What is the benefit of using Apigee to achieve this goal?

- A. It has dashboards that chart dimensions and metrics to report on APIs.
- B. It replicates banking APIs to create new business value.
- C. It measures and tracks their total cost of ownership (TCO).
- D. It allows developers to connect the banking APIs with the public cloud.

Answer: A

Explanation:

Apigee includes analytics services which allow enterprises to report on various aspects of an API.

NEW QUESTION 319

- (Topic 2)

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

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

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

Answer: C

Explanation:

The above is a real world example of a company named Planet, where they sent around 80+ satellites to take pictures of earth every day, 24 hours. They run around 40,000 preemptible VMs concurrently.

Preemptible instances function like normal instances but have the following limitations: Compute Engine might stop preemptible instances at any time due to system events. The probability that Compute Engine will stop a preemptible instance for a system event is generally low, but might vary from day to day and from zone to zone depending on current conditions.

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

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

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

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

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

Important: Spot VMs are the latest version of preemptible VMs. New and existing preemptible VMs continue to be supported, and preemptible VMs use the same pricing model as Spot VMs. However, Spot VMs provide new features that preemptible VMs do not support. For example, preemptible VMs can only run for up to 24 hours at a time, but Spot VMs do not have a maximum runtime. [Learn more about Spot VMs and how to create Spot VMs](#)

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

NEW QUESTION 324

- (Topic 2)

An organization wants to evaluate the performance of their entire cloud infrastructure, including metrics like server uptime and response rate reports. Which Google Cloud tool should the organization use?

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

Answer: B

Explanation:

Because Cloud Monitoring enables users to monitor the performance of their entire cloud infrastructure.

NEW QUESTION 328

- (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 330

- (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 335

- (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 339

- (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 342

- (Topic 1)

An organization has created an ecommerce website. What data on this website would be considered structured data?

- A. Product photographs
- B. Product reviews
- C. Product descriptions
- D. Product ratings score

Answer: D

Explanation:

Because product ratings are structured because they are numerical scores.

NEW QUESTION 345

- (Topic 1)

Your company needs to segment Google Cloud resources used by each team from the others. The teams' efforts are changing frequently, and you need to reduce operational risk and maintain cost visibility. Which approach does Google recommend?

- A. One project per team.
- B. One organization per team.
- C. One project that contains all of each team's resources.
- D. One top-level folder per team.

Answer: A

Explanation:

Reference: <https://cloud.google.com/security/infrastructure/design>
The Teams need to be segmented to have visibility on the resources each team consumes

NEW QUESTION 350

- (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 354

- (Topic 1)

Your ed-tech start-up was originally launched in a small geography. Any user sign-ups, course progress, tests taken, etc. are captured on a self-managed MySQL database. Every user generates many such transactions. Now you're taking the application globally and preparing for a much larger influx of users from all over the world. The existing MySQL server is unlikely to be able to scale. Which convenient option can be considered?

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

Answer: B

Explanation:

Cloud Spanner is a global scale SQL database that scales extremely well. That would be the best choice.

NEW QUESTION 357

- (Topic 1)

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

- A. Elasticity
- B. Availability
- C. Broad Network Access
- D. 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 361

- (Topic 1)

A prospect wants to be able to store and analyze data. Their analysts already know SQL, but are not familiar with other technologies. Which of these databases can the analysts use without additional training?

- A. Cloud SQL, BigQuery, Datastore
- B. Spanner, Cloud SQL, BigQuery
- C. Cloud SQL, Firestore, Datastore
- D. Cloud SQL, Bigtable, BigQuery

Answer: B

Explanation:

Spanner, Cloud SQL, BigQuery

Spanner- Cloud Spanner is a fully managed, mission-critical, relational database service that offers transactional consistency at global scale, automatic, synchronous replication for high availability, and support for two SQL Google Standard SQL and PostgreSQL. Cloud SQL- Cloud SQL is a fully-managed database service that helps you set up,

maintain, manage, and administer your relational databases on Google Cloud Platform. BigQuery- Google BigQuery is a cloud-based Architecture and provides exceptional performance as it can auto-scale up and down based on the data load and performs data analysis efficiently. On the other hand, SQL Server is based on client-server architecture and has fixed performance throughout unless the user scales it manually.

NEW QUESTION 365

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 367

- (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 372

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 373

- (Topic 1)

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

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

Answer: D

Explanation:

Features of Cloud Spanner

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

NEW QUESTION 378

- (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 380

.....

Relate Links

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

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

Contact us

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

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