Google

Exam Questions Professional-Cloud-Architect

Google Certified Professional - Cloud Architect (GCP)

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
- (Exam Topic 1)
For this question, refer to the Mountkirk Games case study.
Mountkirk Games has deployed their new backend on Google Cloud Platform (GCP). You want to create a thorough testing process for new versions of the backend before they are released to the public. You want the testing environment to scale in an economical way. How should you design the process?

A. Create a scalable environment in GCP for simulating production load.
B. Use the existing infrastructure to test the GCP-based backend at scale.
C. Build stress tests into each component of your application using resources internal to GCP to simulate load.
D. Create a set of static environments in GCP to test different levels of load — for example, high, medium, and low.

Answer: A

Explanation:
From scenario: Requirements for Game Backend Platform
≫ Dynamically scale up or down based on game activity
≫ Connect to a managed NoSQL database service
≫ Run customize Linux distro

NEW QUESTION 2
- (Exam Topic 1)
For this question, refer to the Mountkirk Games case study.
Mountkirk Games wants to set up a continuous delivery pipeline. Their architecture includes many small services that they want to be able to update and roll back quickly. Mountkirk Games has the following requirements:
• Services are deployed redundantly across multiple regions in the US and Europe.
• Only frontend services are exposed on the public internet.
• They can provide a single frontend IP for their fleet of services.
• Deployment artifacts are immutable. Which set of products should they use?

A. Google Cloud Storage, Google Cloud Dataflow, Google Compute Engine
B. Google Cloud Storage, Google App Engine, Google Network Load Balancer
C. Google Kubernetes Registry, Google Container Engine, Google HTTP(S) Load Balancer
D. Google Cloud Functions, Google Cloud Pub/Sub, Google Cloud Deployment Manager

Answer: C

NEW QUESTION 3
- (Exam Topic 2)
For this question, refer to the TerramEarth case study.
TerramEarth has equipped unconnected trucks with servers and sensors to collet telemetry data. Next year they want to use the data to train machine learning models. They want to store this data in the cloud while reducing costs. What should they do?

A. Have the vehicle' computer compress the data in hourly snapshots, and store it in a Google Cloud storage (GCS) Nearline bucket.
B. Push the telemetry data in Real-time to a streaming dataflow job that compresses the data, and store it in Google BigQuery.
C. Push the telemetry data in real-time to a streaming dataflow job that compresses the data, and store it in Cloud Bigtable.
D. Have the vehicle's computer compress the data in hourly snapshots, a Store it in a GCS Coldline bucket.

Answer: D

Explanation:
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, costs for data access, and higher per-operation costs. For example:
Cold Data Storage - Infrequently accessed data, such as data stored for legal or regulatory reasons, can be stored at low cost as Coldline Storage, and be available when you need it.
Disaster recovery - In the event of a disaster recovery event, recovery time is key. Cloud Storage provides low latency access to data stored as Coldline Storage.
References: https://cloud.google.com/storage/docs/storage-classes

NEW QUESTION 4
- (Exam Topic 2)
For this question, refer to the TerramEarth case study.
The TerramEarth development team wants to create an API to meet the company's business requirements. You want the development team to focus their development effort on business value versus creating a custom framework. Which method should they use?

A. Use Google App Engine with Google Cloud Endpoint
B. Focus on an API for dealers and partners.
C. Use Google App Engine with a JAX-RS Jersey Java-based framewor
D. Focus on an API for the public.
E. Use Google App Engine with the Swagger (open API Specification) framewor
F. Focus on an API for the public.
G. Use Google Container Engine with a Django Python containe
H. Focus on an API for the public.
I. Use Google Container Engine with a Tomcat container with the Swagger (Open API Specification) framewor
J. Focus on an API for dealers and partners.

Answer: A

Explanation:

https://cloud.google.com/endpoints/docs/openapi/about-cloud-endpoints?hl=en_US&_ga=2.21787131.-1712523
https://cloud.google.com/endpoints/docs/openapi/architecture-overview
https://cloud.google.com/storage/docs/gsutil/commands/test
Develop, deploy, protect and monitor your APIs with Google Cloud Endpoints. Using an Open API Specification or one of our API frameworks, Cloud Endpoints gives you the tools you need for every phase of API development.
From scenario: Business Requirements
Decrease unplanned vehicle downtime to less than 1 week, without increasing the cost of carrying surplus inventory
Support the dealer network with more data on how their customers use their equipment to better position new products and services
Have the ability to partner with different companies – especially with seed and fertilizer suppliers in the fast-growing agricultural business – to create compelling joint offerings for their customers.
Reference: https://cloud.google.com/certification/guides/cloud-architect/casestudy-terramearth

NEW QUESTION 5
- (Exam Topic 5)
You need to ensure reliability for your application and operations by supporting reliable task a scheduling for compute on GCP. Leveraging Google best practices, what should you do?

A. Using the Cron service provided by App Engine, publishing messages directly to a message-processing utility service running on Compute Engine instances.
B. Using the Cron service provided by App Engine, publish messages to a Cloud Pub/Sub topi
C. Subscribe to that topic using a message-processing utility service running on Compute Engine instances.
D. Using the Cron service provided by Google Kubernetes Engine (GKE), publish messages directly to a message-processing utility service running on Compute Engine instances.
E. Using the Cron service provided by GKE, publish messages to a Cloud Pub/Sub topi
F. Subscribe to that topic using a message-processing utility service running on Compute Engine instances.

Answer: B

Explanation:
https://cloud.google.com/solutions/reliable-task-scheduling-compute-engine

NEW QUESTION 6
- (Exam Topic 5)
Your customer is receiving reports that their recently updated Google App Engine application is taking approximately 30 seconds to load for some of their users. This behavior was not reported before the update. What strategy should you take?

A. Work with your ISP to diagnose the problem.
B. Open a support ticket to ask for network capture and flow data to diagnose the problem, then roll back your application.
C. Roll back to an earlier known good release initially, then use Stackdriver Trace and logging to diagnose the problem in a development/test/staging environment.
D. Roll back to an earlier known good release, then push the release again at a quieter period to investigate.Then use Stackdriver Trace and logging to diagnose the problem.

Answer: C

Explanation:
Stackdriver Logging allows you to store, search, analyze, monitor, and alert on log data and events from Google Cloud Platform and Amazon Web Services (AWS). Our API also allows ingestion of any custom log data from any source. Stackdriver Logging is a fully managed service that performs at scale and can ingest application and system log data from thousands of VMs. Even better, you can analyze all that log data in real time.
References: https://cloud.google.com/logging/

NEW QUESTION 7
- (Exam Topic 5)
You are using Cloud CDN to deliver static HTTP(S) website content hosted on a Compute Engine instance group. You want to improve the cache hit ratio. What should you do?

A. Customize the cache keys to omit the protocol from the key.
B. Shorten the expiration time of the cached objects.
C. Make sure the HTTP(S) header "Cache-Region" points to the closest region of your users.
D. Replicate the static content in a Cloud Storage bucke
E. Point CloudCDN toward a load balancer on that bucket.

Answer: A

Explanation:
Reference https://cloud.google.com/cdn/docs/bestpractices#using_custom_cache_keys_to_improve_cache_hit_ratio

NEW QUESTION 8
- (Exam Topic 5)
Your customer runs a web service used by e-commerce sites to offer product recommendations to users. The company has begun experimenting with a machine learning model on Google Cloud Platform to improve the quality of results.
What should the customer do to improve their model's results over time?

A. Export Cloud Machine Learning Engine performance metrics from Stackdriver to BigQuery, to be used to analyze the efficiency of the model.
B. Build a roadmap to move the machine learning model training from Cloud GPUs to Cloud TPUs, which offer better results.
C. Monitor Compute Engine announcements for availability of newer CPU architectures, and deploy the model to them as soon as they are available for additional performance.
D. Save a history of recommendations and results of the recommendations in BigQuery, to be used as training data.

Answer: D

Explanation:
https://cloud.google.com/solutions/building-a-serverless-ml-model

NEW QUESTION 9
- (Exam Topic 5)
You are building a continuous deployment pipeline for a project stored in a Git source repository and want to ensure that code changes can be verified deploying to production. What should you do?

A. Use Spinnaker to deploy builds to production using the red/black deployment strategy so that changes can easily be rolled back.
B. Use Spinnaker to deploy builds to production and run tests on production deployments.
C. Use Jenkins to build the staging branches and the master branc
D. Build and deploy changes to production for 10% of users before doing a complete rollout.
E. Use Jenkins to monitor tags in the repositor
F. Deploy staging tags to a staging environment for testing.After testing, tag the repository for production and deploy that to the production environment.

Answer: D

Explanation:
Reference: https://github.com/GoogleCloudPlatform/continuous-deployment-on-kubernetes/blob/master/ README.md

NEW QUESTION 10
- (Exam Topic 5)
Your company is using BigQuery as its enterprise data warehouse. Data is distributed over several Google Cloud projects. All queries on BigQuery need to be billed on a single project. You want to make sure that no query costs are incurred on the projects that contain the data. Users should be able to query the datasets, but not edit them.
How should you configure users' access roles?

A. Add all users to a grou
B. Grant the group the role of BigQuery user on the billing project and BigQuery dataViewer on the projects that contain the data.
C. Add all users to a grou
D. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery user on the projects that contain the data.
E. Add all users to a grou
F. Grant the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data.
G. Add all users to a grou
H. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery jobUser on the projects that contain the data.

Answer: A

Explanation:
Reference: https://cloud.google.com/bigquery/docs/running-queries

NEW QUESTION 10
- (Exam Topic 5)
The operations manager asks you for a list of recommended practices that she should consider when migrating a J2EE application to the cloud. Which three practices should you recommend? Choose 3 answers

A. Port the application code to run on Google App Engine.
B. Integrate Cloud Dataflow into the application to capture real-time metrics.
C. Instrument the application with a monitoring tool like Stackdriver Debugger.
D. Select an automation framework to reliably provision the cloud infrastructure.
E. Deploy a continuous integration tool with automated testing in a staging environment.
F. Migrate from MySQL to a managed NoSQL database like Google Cloud Datastore or Bigtable.

Answer: AEF

Explanation:
References: https://cloud.google.com/appengine/docs/standard/java/tools/uploadinganapp https://cloud.google.com/appengine/docs/standard/java/building-app/cloud-sql

NEW QUESTION 13
- (Exam Topic 5)
You have an App Engine application that needs to be updated. You want to test the update with production traffic before replacing the current application version. What should you do?

A. Deploy the update using the Instance Group Updater to create a partial rollout, which allows for canary testing.
B. Deploy the update as a new version in the App Engine application, and split traffic between the new and current versions.
C. Deploy the update in a new VPC, and use Google's global HTTP load balancing to split traffic between the update and current applications.
D. Deploy the update as a new App Engine application, and use Google's global HTTP load balancing to split traffic between the new and current applications.

Answer: B

Explanation:
https://cloud.google.com/appengine/docs/standard/python/splitting-traffic

NEW QUESTION 16
- (Exam Topic 5)
Your customer is moving their corporate applications to Google Cloud Platform. The security team wants detailed visibility of all projects in the organization. You

provision the Google Cloud Resource Manager and set up yourself as the org admin. What Google Cloud Identity and Access Management (Cloud IAM) roles should you give to the security team'?

A. Org viewer, project owner
B. Org viewer, project viewer
C. Org admin, project browser
D. Project owner, network admin

Answer: B

Explanation:
https://cloud.google.com/iam/docs/using-iam-securely

NEW QUESTION 19
- (Exam Topic 5)
Your company's test suite is a custom C++ application that runs tests throughout each day on Linux virtual machines. The full test suite takes several hours to complete, running on a limited number of on premises servers reserved for testing. Your company wants to move the testing infrastructure to the cloud, to reduce the amount of time it takes to fully test a change to the system, while changing the tests as little as possible. Which cloud infrastructure should you recommend?

A. Google Compute Engine unmanaged instance groups and Network Load Balancer
B. Google Compute Engine managed instance groups with auto-scaling
C. Google Cloud Dataproc to run Apache Hadoop jobs to process each test
D. Google App Engine with Google Stackdriver for logging

Answer: B

Explanation:
https://cloud.google.com/compute/docs/instance-groups/
Google Compute Engine enables users to launch virtual machines (VMs) on demand. VMs can be launched from the standard images or custom images created by users.
Managed instance groups offer autoscaling capabilities that allow you to automatically add or remove instances from a managed instance group based on increases or decreases in load. Autoscaling helps your applications gracefully handle increases in traffic and reduces cost when the need for resources is lower.

NEW QUESTION 22
- (Exam Topic 5)
A small number of API requests to your microservices-based application take a very long time. You know that each request to the API can traverse many services. You want to know which service takes the longest in those cases. What should you do?

A. Set timeouts on your application so that you can fail requests faster.
B. Send custom metrics for each of your requests to Stackdriver Monitoring.
C. Use Stackdriver Monitoring to look for insights that show when your API latencies are high.
D. Instrument your application with Stackdnver Trace in order to break down the request latencies at each microservice.

Answer: D

Explanation:
https://cloud.google.com/trace/docs/overview

NEW QUESTION 26
- (Exam Topic 5)
A development manager is building a new application He asks you to review his requirements and identify what cloud technologies he can use to meet them. The application must
* 1. Be based on open-source technology for cloud portability
* 2. Dynamically scale compute capacity based on demand
* 3. Support continuous software delivery
* 4. Run multiple segregated copies of the same application stack
* 5. Deploy application bundles using dynamic templates
* 6. Route network traffic to specific services based on URL
Which combination of technologies will meet all of his requirements?

A. Google Container Engine, Jenkins, and Helm
B. Google Container Engine and Cloud Load Balancing
C. Google Compute Engine and Cloud Deployment Manager
D. Google Compute Engine, Jenkins, and Cloud Load Balancing

Answer: A

Explanation:
Helm for managing Kubernetes
Kubernetes can base on the URL to route traffic to different location (path) https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer eg.apiVersion: networking.k8s.io/v1beta1
kind: Ingress metadata:
name: fanout-ingress spec:
rules:
- http: paths:
- path: /* backend: serviceName: web servicePort: 8080
- path: /v2/* backend: serviceName: web2 servicePort: 8080

NEW QUESTION 30

- (Exam Topic 5)
You set up an autoscaling instance group to serve web traffic for an upcoming launch. After configuring the instance group as a backend service to an HTTP(S) load balancer, you notice that virtual machine (VM) instances are being terminated and re-launched every minute. The instances do not have a public IP address. You have verified the appropriate web response is coming from each instance using the curl command. You want to ensure the backend is configured correctly. What should you do?

A. Ensure that a firewall rule exists to allow source traffic on HTTP/HTTPS to reach the load balancer.
B. Assign a public IP to each instance and configure a firewall rule to allow the load balancer to reach the instance public IP.
C. Ensure that a firewall rule exists to allow load balancer health checks to reach the instances in the instance group.
D. Create a tag on each instance with the name of the load balance
E. Configure a firewall rule with the name of the load balancer as the source and the instance tag as the destination.

Answer: C

Explanation:
https://cloud.google.com/vpc/docs/using-firewalls
The best practice when configuration a health check is to check health and serve traffic on the same port. However, it is possible to perform health checks on one port, but serve traffic on another. If you do use two different ports, ensure that firewall rules and services running on instances are configured appropriately. If you run health checks and serve traffic on the same port, but decide to switch ports at some point, be sure to update both the backend service and the health check. Backend services that do not have a valid global forwarding rule referencing it will not be health checked and will have no health status.
References: https://cloud.google.com/compute/docs/load-balancing/http/backend-service

NEW QUESTION 31
- (Exam Topic 5)
Your organization wants to control IAM policies for different departments independently, but centrally. Which approach should you take?

A. Multiple Organizations with multiple Folders
B. Multiple Organizations, one for each department
C. A single Organization with Folder for each department
D. A single Organization with multiple projects, each with a central owner

Answer: C

Explanation:
Folders are nodes in the Cloud Platform Resource Hierarchy. A folder can contain projects, other folders, or a combination of both. You can use folders to group projects under an organization in a hierarchy. For example, your organization might contain multiple departments, each with its own set of GCP resources. Folders allow you to group these resources on a per-department basis. Folders are used to group resources that share common IAM policies. While a folder can contain multiple folders or resources, a given folder or resource can have exactly one parent.
References: https://cloud.google.com/resource-manager/docs/creating-managing-folders

NEW QUESTION 33
- (Exam Topic 5)
Your customer wants to capture multiple GBs of aggregate real-time key performance indicators (KPIs) from their game servers running on Google Cloud Platform and monitor the KPIs with low latency. How should they capture the KPIs?

A. Store time-series data from the game servers in Google Bigtable, and view it using Google Data Studio.
B. Output custom metrics to Stackdriver from the game servers, and create a Dashboard in Stackdriver Monitoring Console to view them.
C. Schedule BigQuery load jobs to ingest analytics files uploaded to Cloud Storage every ten minutes, and visualize the results in Google Data Studio.
D. Insert the KPIs into Cloud Datastore entities, and run ad hoc analysis and visualizations of them inCloud Datala

Answer: A

Explanation:
https://cloud.google.com/monitoring/api/v3/metrics-details#metric-kinds

NEW QUESTION 37
- (Exam Topic 5)
Your marketing department wants to send out a promotional email campaign. The development team wants to minimize direct operation management. They project a wide range of possible customer responses, from 100 to 500,000 click-throughs per day. The link leads to a simple website that explains the promotion and collects user information and preferences. Which infrastructure should you recommend? (CHOOSE TWO)

A. Use Google App Engine to serve the website and Google Cloud Datastore to store user data.
B. Use a Google Container Engine cluster to serve the website and store data to persistent disk.
C. Use a managed instance group to serve the website and Google Cloud Bigtable to store user data.
D. Use a single compute Engine virtual machine (VM) to host a web server, backed by Google Cloud SQL.

Answer: AC

Explanation:
Reference: https://cloud.google.com/storage-options/ References: https://cloud.google.com/storage-options/

NEW QUESTION 42
- (Exam Topic 5)
You are analyzing and defining business processes to support your startup's trial usage of GCP, and you don't yet know what consumer demand for your product will be. Your manager requires you to minimize GCP service costs and adhere to Google best practices. What should you do?

A. Utilize free tier and sustained use discount
B. Provision a staff position for service cost management.
C. Utilize free tier and sustained use discount

D. Provide training to the team about service cost management.
E. Utilize free tier and committed use discount
F. Provision a staff position for service cost management.
G. Utilize free tier and committed use discount
H. Provide training to the team about service cost management.

Answer: D

Explanation:
https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#billing_and_management

NEW QUESTION 47
- (Exam Topic 5)
Your applications will be writing their logs to BigQuery for analysis. Each application should have its own table.
Any logs older than 45 days should be removed. You want to optimize storage and follow Google recommended practices. What should you do?

A. Configure the expiration time for your tables at 45 days
B. Make the tables time-partitioned, and configure the partition expiration at 45 days
C. Rely on BigQuery's default behavior to prune application logs older than 45 days
D. Create a script that uses the BigQuery command line tool (bq) to remove records older than 45 days

Answer: B

Explanation:
https://cloud.google.com/bigquery/docs/managing-partitioned-tables

NEW QUESTION 48
- (Exam Topic 5)
A production database virtual machine on Google Compute Engine has an ext4-formatted persistent disk for data files The database is about to run out of storage space How can you remediate the problem with the least amount of downtime?

A. In the Cloud Platform Console, increase the size of the persistent disk and use the resize2fs command in Linux.
B. Shut down the virtual machine, use the Cloud Platform Console to increase the persistent disk size, then restart the virtual machine.
C. In the Cloud Platform Console, increase the size of the persistent disk and verify the new space is ready to use with the fdisk command in Linux.
D. In the Cloud Platform Console, create a new persistent disk attached to the virtual machine, format and mount it, and configure the database service to move the files to the new disk.
E. In the Cloud Platform Console, create a snapshot of the persistent disk, restore the snapshot to a new larger disk, unmount the old disk, mount the new disk, and restart the database service.

Answer: A

Explanation:
On Linux instances, connect to your instance and manually resize your partitions and file systems to use the additional disk space that you added.
Extend the file system on the disk or the partition to use the added space. If you grew a partition on your disk, specify the partition. If your disk does not have a partition table, specify only the disk ID.
sudo resize2fs /dev/[DISK_ID][PARTITION_NUMBER]
where [DISK_ID] is the device name and [PARTITION_NUMBER] is the partition number for the device where you are resizing the file system.
References: https://cloud.google.com/compute/docs/disks/add-persistent-disk

NEW QUESTION 52
- (Exam Topic 5)
You need to develop procedures to verify resilience of disaster recovery for remote recovery using GCP. Your production environment is hosted on-premises. You need to establish a secure, redundant connection between your on premises network and the GCP network.
What should you do?

A. Verify that Dedicated Interconnect can replicate files to GC
B. Verify that direct peering can establish a secure connection between your networks if Dedicated Interconnect fails.
C. Verify that Dedicated Interconnect can replicate files to GC
D. Verify that Cloud VPN can establish a secure connection between your networks if Dedicated Interconnect fails.
E. Verify that the Transfer Appliance can replicate files to GC
F. Verify that direct peering can establish a secure connection between your networks if the Transfer Appliance fails.
G. Verify that the Transfer Appliance can replicate files to GC
H. Verify that Cloud VPN can establish a secure connection between your networks if the Transfer Appliance fails.

Answer: B

Explanation:
https://cloud.google.com/interconnect/docs/how-to/direct-peering

NEW QUESTION 56
- (Exam Topic 5)
As part of implementing their disaster recovery plan, your company is trying to replicate their production MySQL database from their private data center to their GCP project using a Google Cloud VPN connection. They are experiencing latency issues and a small amount of packet loss that is disrupting the replication. What should they do?

A. Configure their replication to use UDP.
B. Configure a Google Cloud Dedicated Interconnect.
C. Restore their database daily using Google Cloud SQL.
D. Add additional VPN connections and load balance them.

E. Send the replicated transaction to Google Cloud Pub/Sub.

Answer: B

NEW QUESTION 59
- (Exam Topic 5)
You are creating a solution to remove backup files older than 90 days from your backup Cloud Storage bucket. You want to optimize ongoing Cloud Storage spend. What should you do?

A. Write a lifecycle management rule in XML and push it to the bucket with gsutil.
B. Write a lifecycle management rule in JSON and push it to the bucket with gsutil.
C. Schedule a cron script using gsutil is -lr gs://backups/** to find and remove items older than 90 days.
D. Schedule a cron script using gsutil ls -1 gs://backups/** to find and remove items older than 90 days and schedule it with cron.

Answer: B

Explanation:
https://cloud.google.com/storage/docs/gsutil/commands/lifecycle

NEW QUESTION 62
- (Exam Topic 5)
You have an outage in your Compute Engine managed instance group: all instance keep restarting after 5 seconds. You have a health check configured, but autoscaling is disabled. Your colleague, who is a Linux expert, offered to look into the issue. You need to make sure that he can access the VMs. What should you do?

A. Grant your colleague the IAM role of project Viewer
B. Perform a rolling restart on the instance group
C. Disable the health check for the instance grou
D. Add his SSH key to the project-wide SSH keys
E. Disable autoscaling for the instance grou
F. Add his SSH key to the project-wide SSH Keys

Answer: C

Explanation:
https://cloud.google.com/compute/docs/instance-groups/autohealing-instances-in-migs
Health checks used for autohealing should be conservative so they don't preemptively delete and recreate your instances. When an autohealer health check is too aggressive, the autohealer might mistake busy instances for failed instances and unnecessarily restart them, reducing availability

NEW QUESTION 67
- (Exam Topic 5)
You need to evaluate your team readiness for a new GCP project. You must perform the evaluation and create a skills gap plan incorporates the business goal of cost optimization. Your team has deployed two GCP projects successfully to date. What should you do?

A. Allocate budget for team trainin
B. Set a deadline for the new GCP project.
C. Allocate budget for team trainin
D. Create a roadmap for your team to achieve Google Cloud certification based on job role.
E. Allocate budget to hire skilled external consultant
F. Set a deadline for the new GCP project.
G. Allocate budget to hire skilled external consultant
H. Create a roadmap for your team to achieve Google Cloud certification based on job role.

Answer: B

Explanation:
https://services.google.com/fh/files/misc/cloud_center_of_excellence.pdf

NEW QUESTION 68
- (Exam Topic 5)
Google Cloud Platform resources are managed hierarchically using organization, folders, and projects. When Cloud Identity and Access Management (IAM) policies exist at these different levels, what is the effective policy at a particular node of the hierarchy?

A. The effective policy is determined only by the policy set at the node
B. The effective policy is the policy set at the node and restricted by the policies of its ancestors
C. The effective policy is the union of the policy set at the node and policies inherited from its ancestors
D. The effective policy is the intersection of the policy set at the node and policies inherited from its ancestors

Answer: B

Explanation:
Reference: https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy

NEW QUESTION 73
- (Exam Topic 5)
Your company has successfully migrated to the cloud and wants to analyze their data stream to optimize operations. They do not have any existing code for this analysis, so they are exploring all their options. These options include a mix of batch and stream processing, as they are running some hourly jobs and

live-processing some data as it comes in. Which technology should they use for this?

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

Answer: B

Explanation:
Dataflow is for processing both the Batch and Stream.
Cloud Dataflow is a fully-managed service for transforming and enriching data in stream (real time) and batch (historical) modes with equal reliability and expressiveness -- no more complex workarounds or compromises needed.
References: https://cloud.google.com/dataflow/

NEW QUESTION 78
- (Exam Topic 5)
You are designing a mobile chat application. You want to ensure people cannot spoof chat messages, by providing a message were sent by a specific user. What should you do

A. Tag messages client side with the originating user identifier and the destination user.
B. Encrypt the message client side using block-based encryption with a shared key.
C. Use public key infrastructure (PKI) to encrypt the message client side using the originating user's private key.
D. Use a trusted certificate authority to enable SSL connectivity between the client application and the server.

Answer: C

NEW QUESTION 80
- (Exam Topic 5)
You have been engaged by your client to lead the migration of their application infrastructure to GCP. One of their current problems is that the on-premises high performance SAN is requiring frequent and expensive upgrades to keep up with the variety of workloads that are identified as follows: 20TB of log archives retained for legal reasons; 500 GB of VM boot/data volumes and templates; 500 GB of image thumbnails; 200 GB of customer session state data that allows customers to restart sessions even if off-line for several days.
Which of the following best reflects your recommendations for a cost-effective storage allocation?

A. Local SSD for customer session state dat
B. Lifecycle-managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.
C. Memcache backed by Cloud Datastore for the customer session state dat
D. Lifecycle- managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.
E. Memcache backed by Cloud SQL for customer session state dat
F. Assorted local SSD-backed instances for VM boot/data volume
G. Cloud Storage for log archives and thumbnails.
H. Memcache backed by Persistent Disk SSD storage for customer session state dat
I. Assorted local SSDbacked instances for VM boot/data volume
J. Cloud Storage for log archives and thumbnails.

Answer: D

Explanation:
https://cloud.google.com/compute/docs/disks

NEW QUESTION 84
- (Exam Topic 5)
You are creating an App Engine application that uses Cloud Datastore as its persistence layer. You need to retrieve several root entities for which you have the identifiers. You want to minimize the overhead in operations performed by Cloud Datastore. What should you do?

A. Create the Key object for each Entity and run a batch get operation
B. Create the Key object for each Entity and run multiple get operations, one operation for each entity
C. Use the identifiers to create a query filter and run a batch query operation
D. Use the identifiers to create a query filter and run multiple query operations, one operation for each entity

Answer: C

Explanation:
https://cloud.google.com/datastore/docs/concepts/entities#datastore-datastore-batch-upsert-nodejs

NEW QUESTION 88
- (Exam Topic 5)
You need to reduce the number of unplanned rollbacks of erroneous production deployments in your company's web hosting platform. Improvement to the QA/Test processes accomplished an 80% reduction. Which additional two approaches can you take to further reduce the rollbacks? Choose 2 answers

A. Introduce a green-blue deployment model.
B. Replace the QA environment with canary releases.
C. Fragment the monolithic platform into microservices.
D. Reduce the platform's dependency on relational database systems.
E. Replace the platform's relational database systems with a NoSQL database.

Answer: AC

NEW QUESTION 90
- (Exam Topic 5)
An application development team believes their current logging tool will not meet their needs for their new cloud-based product. They want a bettor tool to capture errors and help them analyze their historical log data. You want to help them find a solution that meets their needs, what should you do?

A. Direct them to download and install the Google StackDriver logging agent.
B. Send them a list of online resources about logging best practices.
C. Help them define their requirements and assess viable logging tools.
D. Help them upgrade their current tool to take advantage of any new features.

Answer: C

Explanation:
Help them define their requirements and assess viable logging tools. They know the requirements and the existing tools' problems. While it's true StackDriver Logging and Error Reporting possibly meet all their requirements, there might be other tools also meet their need. They need you to provide expertise to make assessment for new tools, specifically, logging tools that can "capture errors and help them analyze their historical log data".
References: https://cloud.google.com/logging/docs/agent/installation

NEW QUESTION 94
- (Exam Topic 5)
Your company is forecasting a sharp increase in the number and size of Apache Spark and Hadoop jobs being run on your local datacenter You want to utilize the cloud to help you scale this upcoming demand with the least amount of operations work and code change. Which product should you use?

A. Google Cloud Dataflow
B. Google Cloud Dataproc
C. Google Compute Engine
D. Google Container Engine

Answer: B

Explanation:
Google Cloud Dataproc is a fast, easy-to-use, low-cost and fully managed service that lets you run the Apache Spark and Apache Hadoop ecosystem on Google Cloud Platform. Cloud Dataproc provisions big or small clusters rapidly, supports many popular job types, and is integrated with other Google Cloud Platform services, such as Google Cloud Storage and Stackdriver Logging, thus helping you reduce TCO.
References: https://cloud.google.com/dataproc/docs/resources/faq

NEW QUESTION 95
- (Exam Topic 5)
You have a Python web application with many dependencies that requires 0.1 CPU cores and 128 MB of memory to operate in production. You want to monitor and maximize machine utilization. You also to reliably deploy new versions of the application. Which set of steps should you take?

A. Perform the following:1) Create a managed instance group with f1-micro type machines.2) Use a startup script to clone the repository, check out the production branch, install the dependencies, and start the Python app.3) Restart the instances to automatically deploy new production releases.
B. Perform the following:1) Create a managed instance group with n1-standard-1 type machines.2) Build a Compute Engine image from the production branch that contains all of the dependencies and automatically starts the Python app.3) Rebuild the Compute Engine image, and update the instance template to deploy new production releases.
C. Perform the following:1) Create a Kubernetes Engine cluster with n1-standard-1 type machines.2) Build a Docker image from the production branch with all of the dependencies, and tag it with the version number.3) Create a Kubernetes Deployment with the imagePullPolicy set to "IfNotPresent" in the staging namespace, and then promote it to the production namespace after testing.
D. Perform the following:1) Create a Kubernetes Engine cluster with n1-standard-4 type machines.2) Build a Docker image from the master branch will all of the dependencies, and tag it with "latest".3) Create a Kubernetes Deployment in the default namespace with the imagePullPolicy set to "Always". Restart the pods to automatically deploy new production releases.

Answer: D

Explanation:
https://cloud.google.com/compute/docs/instance-templates

NEW QUESTION 96
- (Exam Topic 5)
Your company operates nationally and plans to use GCP for multiple batch workloads, including some that are not time-critical. You also need to use GCP services that are HIPAA-certified and manage service costs.
How should you design to meet Google best practices?

A. Provisioning preemptible VMs to reduce cos
B. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.
C. Provisioning preemptible VMs to reduce cos
D. Disable and then discontinue use of all GCP and APIs that are not HIPAA-compliant.
E. Provision standard VMs in the same region to reduce cos
F. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.
G. Provision standard VMs to the same region to reduce cos
H. Disable and then discontinue use of all GCP services and APIs that are not HIPAA-compliant.

Answer: B

Explanation:
https://cloud.google.com/security/compliance/hipaa/

NEW QUESTION 98

- (Exam Topic 5)
You are deploying an application on App Engine that needs to integrate with an on-premises database. For security purposes, your on-premises database must not be accessible through the public Internet. What should you do?

A. Deploy your application on App Engine standard environment and use App Engine firewall rules to limit access to the open on-premises database.
B. Deploy your application on App Engine standard environment and use Cloud VPN to limit access to the onpremises database.
C. Deploy your application on App Engine flexible environment and use App Engine firewall rules to limit access to the on-premises database.
D. Deploy your application on App Engine flexible environment and use Cloud VPN to limit access to the on-premises database.

Answer: D

Explanation:
https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databases

NEW QUESTION 101
- (Exam Topic 5)
You have created several preemptible Linux virtual machine instances using Google Compute Engine. You want to properly shut down your application before the virtual machines are preempted. What should you do?

A. Create a shutdown script named k99.shutdown in the /etc/rc.6.d/ directory.
B. Create a shutdown script registered as a xinetd service in Linux and configure a Stackdnver endpoint check to call the service.
C. Create a shutdown script and use it as the value for a new metadata entry with the key shutdown-script in the Cloud Platform Console when you create the new virtual machine instance.
D. Create a shutdown script, registered as a xinetd service in Linux, and use the gcloud compute instances add-metadata command to specify the service URL as the value for a new metadata entry with the key shutdown-script-url

Answer: C

NEW QUESTION 103
- (Exam Topic 5)
Your company pushes batches of sensitive transaction data from its application server VMs to Cloud Pub/Sub for processing and storage. What is the Google-recommended way for your application to authenticate to the required Google Cloud services?

A. Ensure that VM service accounts are granted the appropriate Cloud Pub/Sub IAM roles.
B. Ensure that VM service accounts do not have access to Cloud Pub/Sub, and use VM access scopes to grant the appropriate Cloud Pub/Sub IAM roles.
C. Generate an OAuth2 access token for accessing Cloud Pub/Sub, encrypt it, and store it in Cloud Storage for access from each VM.
D. Create a gateway to Cloud Pub/Sub using a Cloud Function, and grant the Cloud Function service account the appropriate Cloud Pub/Sub IAM roles.

Answer: A

NEW QUESTION 104
- (Exam Topic 5)
You want to optimize the performance of an accurate, real-time, weather-charting application. The data comes from 50,000 sensors sending 10 readings a second, in the format of a timestamp and sensor reading. Where
should you store the data?

A. Google BigQuery
B. Google Cloud SQL
C. Google Cloud Bigtable
D. Google Cloud Storage

Answer: C

Explanation:
It is time-series data, So Big Table. https://cloud.google.com/bigtable/docs/schema-design-time-series
Google Cloud Bigtable is a scalable, fully-managed NoSQL wide-column database that is suitable for both real-time access and analytics workloads.
Good for:

➢ Low-latency read/write access
➢ High-throughput analytics
➢ Native time series support
➢ Common workloads:
➢ IoT, finance, adtech
➢ Personalization, recommendations
➢ Monitoring
➢ Geospatial datasets
➢ Graphs
References: https://cloud.google.com/storage-options/

NEW QUESTION 109
- (Exam Topic 5)
You have deployed an application to Kubernetes Engine, and are using the Cloud SQL proxy container to make the Cloud SQL database available to the services running on Kubernetes. You are notified that the application is reporting database connection issues. Your company policies require a post-mortem. What should you do?

A. Use gcloud sql instances restart.
B. Validate that the Service Account used by the Cloud SQL proxy container still has the Cloud Build Editor role.

C. In the GCP Console, navigate to Stackdriver Loggin
D. Consult logs for Kubernetes Engine and Cloud SQL.
E. In the GCP Console, navigate to Cloud SQ
F. Restore the latest backu
G. Use kubect1 to restart all pods.

Answer: C

NEW QUESTION 112
- (Exam Topic 5)
Your customer wants to do resilience testing of their authentication layer. This consists of a regional managed instance group serving a public REST API that reads from and writes to a Cloud SQL instance.
What should you do?

A. Engage with a security company to run web scrapes that look your users' authentication data om malicious websites and notify you if any if found.
B. Deploy intrusion detection software to your virtual machines to detect and log unauthorized access.
C. Schedule a disaster simulation exercise during which you can shut off all VMs in a zone to see how your application behaves.
D. Configure a red replica for your Cloud SQL instance in a different zone than the master, and then manually trigger a failover while monitoring KPIs for our REST API.

Answer: C

NEW QUESTION 116
- (Exam Topic 5)
You are deploying a PHP App Engine Standard service with SQL as the backend. You want to minimize the number of queries to the database.
What should you do?

A. Set the memcache service level to dedicate
B. Create a key from the hash of the query, and return database values from memcache before issuing a query to Cloud SQL.
C. Set the memcache service level to dedicate
D. Create a cron task that runs every minute to populate the cache with keys containing query results.
E. Set the memcache service level to share
F. Create a cron task that runs every minute to save all expected queries to a key called "cached-queries".
G. Set the memcache service level to share
H. Create a key called "cached-queries", and return database values from the key before using a query to Cloud SQL.

Answer: A

Explanation:
https://cloud.google.com/appengine/docs/standard/php/memcache/using

NEW QUESTION 119
- (Exam Topic 5)
Your company wants to start using Google Cloud resources but wants to retain their on-premises Active Directory domain controller for identity management. What should you do?

A. Use the Admin Directory API to authenticate against the Active Directory domain controller.
B. Use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML SSO.
C. Use Cloud Identity-Aware Proxy configured to use the on-premises Active Directory domain controller as an identity provider.
D. Use Compute Engine to create an Active Directory (AD) domain controller that is a replica of the onpremises AD domain controller using Google Cloud Directory Sync.

Answer: B

Explanation:
https://cloud.google.com/solutions/federating-gcp-with-active-directory-introduction#implementing_federation

NEW QUESTION 123
- (Exam Topic 5)
You are designing an application for use only during business hours. For the minimum viable product release, you'd like to use a managed product that automatically "scales to zero" so you don't incur costs when there is no activity.
Which primary compute resource should you choose?

A. Cloud Functions
B. Compute Engine
C. Kubernetes Engine
D. AppEngine flexible environment

Answer: A

Explanation:
https://cloud.google.com/serverless-options

NEW QUESTION 128
- (Exam Topic 5)
You are using a single Cloud SQL instance to serve your application from a specific zone. You want to introduce high availability. What should you do?

A. Create a read replica instance in a different region

B. Create a failover replica instance in a different region
C. Create a read replica instance in the same region, but in a different zone
D. Create a failover replica instance in the same region, but in a different zone

Answer: B

Explanation:
https://cloud.google.com/sql/docs/mysql/high-availability

NEW QUESTION 131
- (Exam Topic 5)
Your web application must comply with the requirements of the European Union's General Data Protection Regulation (GDPR). You are responsible for the technical architecture of your web application. What should you do?

A. Ensure that your web application only uses native features and services of Google Cloud Platform,because Google already has various certifications and provides "pass-on" compliance when you use native features.
B. Enable the relevant GDPR compliance setting within the GCPConsole for each of the services in use within your application.
C. Ensure that Cloud Security Scanner is part of your test planning strategy in order to pick up any compliance gaps.
D. Define a design for the security of data in your web application that meets GDPR requirements.

Answer: D

Explanation:
https://cloud.google.com/security/gdpr/?tab=tab4
Reference: https://www.mobiloud.com/blog/gdpr-compliant-mobile-app/

NEW QUESTION 132
- (Exam Topic 5)
Your company is building a new architecture to support its data-centric business focus. You are responsible for setting up the network. Your company's mobile and web-facing applications will be deployed on-premises, and all data analysis will be conducted in GCP. The plan is to process and load 7 years of archived .csv files totaling 900 TB of data and then continue loading 10 TB of data daily. You currently have an existing 100-MB internet connection.
What actions will meet your company's needs?

A. Compress and upload both achieved files and files uploaded daily using the qsutil –m option.
B. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transfer archived data to Cloud Storag
C. Establish a connection with Google using a Dedicated Interconnect or Direct Peering connection and use it to upload files daily.
D. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transferarchived data to Cloud Storag
E. Establish one Cloud VPN Tunnel to VPC networks over the public internet, and compares and upload files daily using the gsutil –m option.
F. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storag
G. Establish a Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily.

Answer: B

Explanation:
https://cloud.google.com/interconnect/docs/how-to/direct-peering

NEW QUESTION 133
- (Exam Topic 5)
You are migrating your on-premises solution to Google Cloud in several phases. You will use Cloud VPN to maintain a connection between your on-premises systems and Google Cloud until the migration is completed.
You want to make sure all your on-premises systems remain reachable during this period. How should you organize your networking in Google Cloud?

A. Use the same IP range on Google Cloud as you use on-premises
B. Use the same IP range on Google Cloud as you use on-premises for your primary IP range and use a secondary range that does not overlap with the range you use on-premises
C. Use an IP range on Google Cloud that does not overlap with the range you use on-premises
D. Use an IP range on Google Cloud that does not overlap with the range you use on-premises for your primary IP range and use a secondary range with the same IP range as you use on-premises

Answer: C

NEW QUESTION 138
......

Relate Links

100% Pass Your Professional-Cloud-Architect Exam with Exambible Prep Materials

https://www.exambible.com/Professional-Cloud-Architect-exam/

Contact us

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

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