Professional-Cloud-Architect Dumps

Google Certified Professional - Cloud Architect (GCP)

https://www.certleader.com/Professional-Cloud-Architect-dumps.html

NEW QUESTION 1
- (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 2
- (Topic 1)
For this question, refer to the Mountkirk Games case study.
Mountkirk Games wants to set up a real-time analytics platform for their new game. The new platform must meet their technical requirements. Which combination of Google technologies will meet all of their requirements?

A. Container Engine, Cloud Pub/Sub, and Cloud SQL
B. Cloud Dataflow, Cloud Storage, Cloud Pub/Sub, and BigQuery
C. Cloud SQL, Cloud Storage, Cloud Pub/Sub, and Cloud Dataflow
D. Cloud Dataproc, Cloud Pub/Sub, Cloud SQL, and Cloud Dataflow
E. Cloud Pub/Sub, Compute Engine, Cloud Storage, and Cloud Dataproc

Answer: B

Explanation:
A real time requires Stream / Messaging so Pub/Sub, Analytics by Big Query.
Ingest millions of streaming events per second from anywhere in the world with Cloud Pub/Sub, powered by Google's unique, high-speed private network. Process the streams with Cloud Dataflow to ensure reliable, exactly-once, low-latency data transformation. Stream the transformed data into BigQuery, the cloud-native data warehousing service, for immediate analysis via SQL or popular visualization tools.
From scenario: They plan to deploy the game's backend on Google Compute Engine so they can capture streaming metrics, run intensive analytics.
Requirements for Game Analytics Platform
? Dynamically scale up or down based on game activity
? Process incoming data on the fly directly from the game servers
? Process data that arrives late because of slow mobile networks
? Allow SQL queries to access at least 10 TB of historical data
? Process files that are regularly uploaded by users' mobile devices
? Use only fully managed services
References: https://cloud.google.com/solutions/big-data/stream-analytics/

NEW QUESTION 3
- (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 4
- (Topic 2)
For this question, refer to the TerramEarth case study.
To speed up data retrieval, more vehicles will be upgraded to cellular connections and be able to transmit data to the ETL process. The current FTP process is error-prone and restarts the data transfer from the start of the file when connections fail, which happens often. You want to improve the reliability of the solution and minimize data transfer time on the cellular connections. What should you do?

A. Use one Google Container Engine cluster of FTP server
B. Save the data to a Multi- Regional bucke
C. Run the ETL process using data in the bucket.
D. Use multiple Google Container Engine clusters running FTP servers located in different region
E. Save the data to Multi-Regional buckets in us, eu, and asi
F. Run the ETL process using the data in the bucket.

G. Directly transfer the files to different Google Cloud Multi-Regional Storage bucket locations in us, eu, and asia using Google APIs over HTTP(S). Run the ETL process using the data in the bucket.
H. Directly transfer the files to a different Google Cloud Regional Storage bucket location in us, eu, and asia using Google APIs over HTTP(S). Run the ETL process to retrieve the data from each Regional bucket.

Answer: D

Explanation:
 https://cloud.google.com/storage/docs/locations

NEW QUESTION 5
- (Topic 3)
For this question, refer to the JencoMart case study.
The JencoMart security team requires that all Google Cloud Platform infrastructure is deployed using a least privilege model with separation of duties for administration between production and development resources. What Google domain and project structure should you recommend?

A. Create two G Suite accounts to manage users: one for development/test/staging andone for productio
B. Each account should contain one project for every application.
C. Create two G Suite accounts to manage users: one with a single project for all development applications and one with a single project for all production applications.
D. Create a single G Suite account to manage users with each stage of each application in its own project.
E. Create a single G Suite account to manage users with one project for the development/test/staging environment and one project for the production environment.

Answer: D

Explanation:
Note: The principle of least privilege and separation of duties are concepts that, although semantically different, are intrinsically related from the standpoint of security. The intent behind both is to prevent people from having higher privilege levels than they actually need
? Principle of Least Privilege: Users should only have the least amount of privileges required to perform their job and no more. This reduces authorization exploitation by limiting access to resources such as targets, jobs, or monitoring templates for which they are not authorized.
? Separation of Duties: Beyond limiting user privilege level, you also limit user duties, or the specific jobs they can perform. No user should be given responsibility for more than one related function. This limits the ability of a user to perform a malicious action and then cover up that action.
References: https://cloud.google.com/kms/docs/separation-of-duties

NEW QUESTION 6
- (Topic 4)
For this question, refer to the Dress4Win case study.
Dress4Win would like to become familiar with deploying applications to the cloud by successfully deploying some applications quickly, as is. They have asked for your recommendation. What should you advise?

A. Identify self-contained applications with external dependencies as a first move to the cloud.
B. Identify enterprise applications with internal dependencies and recommend these as a first move to the cloud.
C. Suggest moving their in-house databases to the cloud and continue serving requests to on-premise applications.
D. Recommend moving their message queuing servers to the cloud and continue handling requests to on-premise applications.

Answer: A

Explanation:
 https://cloud.google.com/blog/products/gcp/the-five-phases-of-migrating-to-google-cloud-platform

NEW QUESTION 7
- (Topic 4)
For this question, refer to the Dress4Win case study.
The Dress4Win security team has disabled external SSH access into production virtual machines (VMs) on Google Cloud Platform (GCP). The operations team needs to remotely manage the VMs, build and push Docker containers, and manage Google Cloud Storage objects. What can they do?

A. Grant the operations engineers access to use Google Cloud Shell.
B. Configure a VPN connection to GCP to allow SSH access to the cloud VMs.
C. Develop a new access request process that grants temporary SSH access to cloud VMs when an operations engineer needs to perform a task.
D. Have the development team build an API service that allows the operations team to execute specific remote procedure calls to accomplish their tasks.

Answer: A

NEW QUESTION 8
- (Topic 4)
For this question, refer to the Dress4Win case study.
As part of Dress4Win's plans to migrate to the cloud, they want to be able to set up a managed logging and monitoring system so they can handle spikes in their traffic load.
They want to ensure that:
• The infrastructure can be notified when it needs to scale up and down to handle the ebb and flow of usage throughout the day
• Their administrators are notified automatically when their application reports errors.
• They can filter their aggregated logs down in order to debug one piece of the application across many hosts
Which Google StackDriver features should they use?

A. Logging, Alerts, Insights, Debug
B. Monitoring, Trace, Debug, Logging
C. Monitoring, Logging, Alerts, Error Reporting
D. Monitoring, Logging, Debug, Error Report

Answer: D

NEW QUESTION 9
- (Topic 4)
The current Dress4win system architecture has high latency to some customers because it is located in one
data center.
As of a future evaluation and optimizing for performance in the cloud, Dresss4win wants to distribute it's system
architecture to multiple locations when Google cloud platform. Which approach should they use?

A. Use regional managed instance groups and a global load balancer to increase performance because theregional managed instance group can grow instances in each region separately based on traffic.
B. Use a global load balancer with a set of virtual machines that forward the requests to a closer group ofvirtual machines managed by your operations team.
C. Use regional managed instance groups and a global load balancer to increase reliability by providingautomatic failover between zones in different regions.
D. Use a global load balancer with a set of virtual machines that forward the requests to a closer group ofvirtual machines as part of a separate managed instance groups.

Answer: A

NEW QUESTION 10
- (Topic 4)
For this question, refer to the Dress4Win case study.
Dress4Win has asked you for advice on how to migrate their on-premises MySQL deployment to the cloud. They want to minimize downtime and performance impact to their on-premises solution during the migration. Which approach should you recommend?

A. Create a dump of the on-premises MySQL master server, and then shut it down, upload it to the cloud environment, and load into a new MySQL cluster.
B. Setup a MySQL replica server/slave in the cloud environment, and configure it for asynchronous replication from the MySQL master server on-premises until cutover.
C. Create a new MySQL cluster in the cloud, configure applications to begin writing to both on-premises and cloud MySQL masters, and destroy the original cluster at cutover.
D. Create a dump of the MySQL replica server into the cloud environment, load it into: Google Cloud Datastore, and configure applications to read/write to Cloud Datastore at cutover.

Answer: B

NEW QUESTION 10
- (Topic 5)
You are responsible for the Google Cloud environment in your company Multiple departments need access to their own projects and the members within each department will have the same project responsibilities You want to structure your Google Cloud environment for minimal maintenance and maximum overview of 1AM permissions as each department's projects start and end You want to follow Google-recommended practices What should you do?

A. Create a Google Group per department and add all department members to their respective groups Create a folder per departmentand grant the respective group the required 1AM permissions at the folder level Add the projects under the respective folders
B. Grant all department members the required 1AM permissions for their respective projects
C. Create a Google Group per department and add all department members to theirrespective groups Grant each group the required I AM permissions for their respective projects
D. Create a folder per department and grant the respective members of the department the required 1AM permissions at the folder leve
E. Structure all projects for each department under the respective folders

Answer: A

Explanation:
 This option follows the Google-recommended practices for structuring a Google Cloud environment for minimal maintenance and maximum overview of IAM permissions. By creating a Google Group per department and adding all department members to their respective groups, you can simplify user management and avoid granting IAM permissions to individual users. By creating a folder per department and granting the respective group the required IAM permissions at the folder level, you can enforce consistent policies across all projects within each department and avoid granting IAM permissions at the project level. By adding the projects under the respective folders, you can organize your resources hierarchically and leverage inheritance of IAM policies from folders to projects. The other options are not optimal for this scenario, because they either require granting IAM permissions to individual users (B, C), or do not use Google Groups to manage users (D). References:
? https://cloud.google.com/architecture/framework/system-design
? https://cloud.google.com/architecture/identity/best-practices-for-planning
? https://cloud.google.com/resource-manager/docs/creating-managing-folders

NEW QUESTION 15
- (Topic 5)
Your company is running its application workloads on Compute Engine. The applications have been deployed in production, acceptance, and development environments. The production environment is business-critical and is used 24/7, while the acceptance and development environments are only critical during office hours. Your CFO has asked you to optimize these environments to achieve cost savings during idle times. What should you do?

A. Create a shell script that uses the gcloud command to change the machine type of the development and acceptance instances to a smaller machine type outside of office hour
B. Schedule the shell script on one of the production instances to automate the task.
C. Use Cloud Scheduler to trigger a Cloud Function that will stop the development and acceptance environments after office hours and start them just before office hours.
D. Deploy the development and acceptance applications on a managed instance group and enable autoscaling.
E. Use regular Compute Engine instances for the production environment, and usepreemptible VMs for the acceptance and development environments.

Answer: B

Explanation:

Reference: https://cloud.google.com/blog/products/it-ops/best-practices-for-optimizing- your-cloud-costs

NEW QUESTION 16
- (Topic 5)
Your company has a networking team and a development team. The development team runs applications on Compute Engine instances that contain sensitive data. The development team requires administrative permissions for Compute Engine. Your company requires all network resources to be managed by the networking team. The development team does not want the networking team to have access to the sensitive data on the instances. What should you do?

A. * 1. Create a project with a standalone VPC and assign the Network Admin role to the networking team.* 2. Create a second project with a standalone VPC and assign the Compute Admin role to the development team.* 3. Use Cloud VPN to join the two VPCs.
B. * 1. Create a project with a standalone Virtual Private Cloud (VPC), assign the Network Admin role to the networking team, and assign the Compute Admin role to the development team.
C. * 1. Create a project with a Shared VPC and assign the Network Admin role to the networking team.* 2. Create a second project without a VPC, configure it as a Shared VPC service project,and assign the Compute Admin role to the development team.
D. * 1. Create a project with a standalone VPC and assign the Network Admin role to the networking team.* 2. Create a second project with a standalone VPC and assign the Compute Admin role to the development team.* 3. Use VPC Peering to join the two VPCs.

Answer: C

Explanation:
In this scenario, a large organization has a central team that manages security and networking controls for the entire organization. Developers do not have permissions to make changes to any network or security settings defined by the security and networking team but they are granted permission to create resources such as virtual machines in shared subnets. To facilitate this the organization makes use of a shared VPC (Virtual Private Cloud). A shared VPC allows creation of a VPC network of RFC 1918 IP spaces that associated projects (service projects) can then use. Developers using the associated projects can create VM instances in the shared VPC network spaces. The organization's network and security admins can create subnets, VPNs, and firewall rules usable by all the projects in the VPC network. https://cloud.google.com/iam/docs/job-functions/networking#single_team_manages_security_network_for_organization
Reference: https://cloud.google.com/vpc/docs/shared-vpc

NEW QUESTION 21
- (Topic 5)
You have developed an application using Cloud ML Engine that recognizes famous paintings from uploaded images. You want to test the application and allow specific people to upload images for the next 24 hours. Not all users have a Google Account. How should you have users upload images?

A. Have users upload the images to Cloud Storag
B. Protect the bucket with a passwordthat expires after 24 hours.
C. Have users upload the images to Cloud Storage using a signed URL that expires after 24 hours.
D. Create an App Engine web application where users can upload image
E. Configure App Engine to disable the application after 24 hour
F. Authenticate users via Cloud Identity.
G. Create an App Engine web application where users can upload images for the next 24 hour
H. Authenticate users via Cloud Identity.

Answer: A

Explanation:
https://cloud.google.com/blog/products/storage-data-transfer/uploading-images-directly-to-cloud-storage-by-using-signed-url

NEW QUESTION 22
- (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 amessage-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 24
- (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 27
- (Topic 5)
Your company has just recently activated Cloud Identity to manage users. The Google Cloud Organization has been configured as wed. The security learn needs to secure protects that will be part of the Organization. They want to prohibit IAM users outside the domain from gaining permissions from now on. What should they do?

A. Configure an organization policy to restrict identities by domain
B. Configure an organization policy to block creation of service accounts
C. Configure Cloud Scheduler o trigger a Cloud Function every hour that removes all users that don't belong to the Cloud identity domain from all projects.
D. Create a technical user (e g . crawler@yourdomain com), and give it the protect owner rote at root organization level Write a bash script that• Lists all me IAM rules of all projects within the organization• Deletes all users that do not belong to the company domainCreate a Compute Engine instance m a project within the Organization and configure gcloud to be executed with technical user credentials Configure a cron job that executes the bash script every hour.

Answer: A

Explanation:
 https://cloud.google.com/resource-manager/docs/organization- policy/restricting-domains

NEW QUESTION 29
- (Topic 5)
Your company is planning to perform a lift and shift migration of their Linux RHEL 6.5+ virtual machines. The virtual machines are running in an on-premises VMware environment. You want to migrate them to Compute Engine following Google- recommended practices. What should you do?

A. * 1. Define a migration plan based on the list of the applications and their dependencies.* 2. Migrate all virtual machines into Compute Engine individually with Migrate for Compute Engine.
B. * 1. Perform an assessment of virtual machines running in the current VMware environment.* 2. Create images of all disk
C. Import disks on Compute Engine.* 3. Create standard virtual machines where the boot disks are the ones you have imported.
D. * 1. Perform an assessment of virtual machines running in the current VMware environment.* 2. Define a migration plan, prepare a Migrate for Compute Engine migration RunBook, and execute the migration.
E. * 1. Perform an assessment of virtual machines running in the current VMware environment.* 2.Install a third-party agent on all selected virtual machine
F. 3.Migrate all virtual machines into Compute Engine.

Answer: C

Explanation:
The framework illustrated in the preceding diagram has four phases:
•Assess. In this phase, you assess your source environment, assess the workloads that you want to migrate to Google Cloud, and assess which VMs support each workload.
•Plan. In this phase, you create the basic infrastructure for Migrate for Compute Engine, such as provisioning the resource hierarchy and setting up network access.
•Deploy. In this phase, you migrate the VMs from the source environment to Compute Engine.
•Optimize. In this phase, you begin to take advantage of the cloud technologies and capabilities.
Reference: https://cloud.google.com/architecture/migrating-vms-migrate-for-compute-engine-getting-started

NEW QUESTION 30
- (Topic 5)
You are developing your microservices application on Google Kubernetes Engine. During testing, you want to validate the behavior of your application in case a specific microservice
should suddenly crash. What should you do?

A. Add a taint to one of the nodes of the Kubernetes cluste
B. For the specific microservice, configure a pod anti-affinity label that has the name of the tainted node as a value.
C. Use Istio's fault injection on the particular microservice whose faulty behavior you want to simulate.
D. Destroy one of the nodes of the Kubernetes cluster to observe the behavior.
E. Configure Istio's traffic management features to steer the traffic away from a crashing microservice.

Answer: B

Explanation:
 Microservice runs on all nodes. The Micro service runs on Pod, Pod runs on Nodes. Nodes is nothing but Virtual machines. Once deployed the application microservices will get deployed across all Nodes. Destroying one node may not mimic the behaviour of microservice crashing as it may be running in other nodes. link: https://istio.io/latest/docs/tasks/traffic-management/fault-injection/

NEW QUESTION 32
- (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 returndatabase 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 databasevalues 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 35
- (Topic 5)
You are moving an application that uses MySQL from on-premises to Google Cloud. The application will run on Compute Engine and will use Cloud SQL. You want to cut over to the Compute Engine deployment of the application with minimal downtime and no data loss to your customers. You want to migrate the application with minimal modification. You also need to determine the cutover strategy. What should you do?

A. * 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server.* 2. Stop the on-premises application.* 3. Create a mysqldump of the on-premises MySQL serve
B. * 4.Upload the dump to a Cloud Storage bucket.* 5. Import the dump into Cloud SQL.* 6. Modify the source code of the application to write queries to both databases and read from its local database.* 7. Start the Compute Engine applicatio
C. * 8.Stop the on-premises application.
D. * 1. Set up Cloud SQL proxy and MySQL prox
E. * 2.Create a mysqldump of the on-premises MySQL serve
F. * 3.Upload the dump to a Cloud Storage bucket.* 4.Import the dump into Cloud SQ
G. * 5.Stop the on-premises applicatio
H. * 6.Start the Compute Engine application.
I. * 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server.* 2. Stop the on-premises application.* 3. Start the Compute Engine application, configured to read and write to the on-premises MySQL server.* 4. Create the replication configuration in Cloud SQL.* 5. Configure the source database server to accept connections from the Cloud SQL replic
J. * 6.Finalize the Cloud SQL replica configuration.* 7.When replication has been completed, stop the Compute Engine applicatio
K. * 8.Promote the Cloud SQL replica to a standalone instance.* 9.Restart the Compute Engine application, configured to read and write to the Cloud SQL standalone instance.
L. * 1. Stop the on-premises application.* 2.Create a mysqldump of the on-premises MySQL serve
M. * 3.Upload the dump to a Cloud Storage bucket.* 4. Import the dump into Cloud SQL.* 5. Start the application on Compute Engine.

Answer: C

Explanation:
External replica promotion migration In the migration strategy of external replica promotion, you create an external database replica and synchronize the existing data to that replica. This can happen with minimal downtime to the existing database. When you have a replica database, the two databases have different roles that are referred to in this document as primary and replica. After the data is synchronized, you promote the replica to be the primary in order to move the management layer with minimal impact to database uptime. In Cloud SQL, an easy way to accomplish the external replica promotion is to use the automated migration workflow. This process automates many of the steps that are needed for this type of migration.
https://cloud.google.com/architecture/migrating-mysql-to-cloudsql-concept
- The best option for migrating your MySQL database is to use an external replica promotion. In this strategy, you create a replica database and set your existing database as the primary. You wait until the two databases are in sync, and you then promote your MySQL replica database to be the primary. This process minimizes database downtime related to the database migration. - https://cloud.google.com/architecture/migrating-mysql-to-cloudsql-concept#external_replica_promotion_migration

NEW QUESTION 40
- (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 45
- (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 reposit or
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 50
- (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 55
- (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 58
- (Topic 5)
To reduce costs, the Director of Engineering has required all developers to move their development infrastructure resources from on-premises virtual machines (VMs) to Google Cloud Platform. These resources go through multiple start/stop events during the day and require state to persist. You have been asked to design the process of running a development environment in Google Cloud while providing cost visibility to the finance department. Which two steps should you take? Choose 2 answers

A. Use the --no-auto-delete flag on all persistent disks and stop the VM.
B. Use the -auto-delete flag on all persistent disks and terminate the VM.
C. Apply VM CPU utilization label and include it in the BigQuery billing export.
D. Use Google BigQuery billing export and labels to associate cost to groups.
E. Store all state into local SSD, snapshot the persistent disks, and terminate the VM.
F. Store all state in Google Cloud Storage, snapshot the persistent disks, and terminate the VM.

Answer: AD

Explanation:
https://cloud.google.com/billing/docs/how-to/export-data-bigquery

NEW QUESTION 60
- (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 63
- (Topic 5)
You are implementing a single Cloud SQL MySQL second-generation database that contains business-critical transaction data. You want to ensure that the

minimum amount of data is lost in case of catastrophic failure. Which two features should you implement? (Choose two.)

A. Sharding
B. Read replicas
C. Binary logging
D. Automated backups
E. Semisynchronous replication

Answer: CD

Explanation:
Backups help you restore lost data to your Cloud SQL instance. Additionally, if an instance is having a problem, you can restore it to a previous state by using the backup to overwrite it. Enable automated backups for any instance that contains necessary data. Backups protect your data from loss or damage.
Enabling automated backups, along with binary logging, is also required for some operations, such as clone and replica creation.
Reference: https://cloud.google.com/sql/docs/mysql/backup-recovery/backups

NEW QUESTION 65
- (Topic 5)
You want to allow your operations learn to store togs from all the production protects in your Organization, without during logs from other projects All of the production projects are contained in a folder. You want to ensure that all logs for existing and new production projects are captured automatically. What should you do?

A. Create an aggregated export on the Production folde
B. Set the log sink to be a Cloud Storage bucket in an operations project
C. Create an aggregated export on the Organization resourc
D. Set the tog sink to be a Cloud Storage bucket in an operations project.
E. Create log exports in the production project
F. Set the log sinks to be a Cloud Storagebucket in an operations project.
G. Create tog exports in the production project
H. Set the tog sinks to be BigQuery datasets in the production projects and grant IAM access to the operations team to run queries on the datasets

Answer: A

Explanation:
? An aggregated export is a type of sink that combines and routes log entries from the Google Cloud resources contained by an organization or folder1. By creating an aggregated export on the Production folder, you can capture all the logs from the existing and new production projects in that folder automatically1.
? A log sink is a destination for log entries that match a filter1. By setting the log sink to be a Cloud Storage bucket in an operations project, you can store the log entries in Cloud Storage and allow your operations team to access them1.

NEW QUESTION 70
- (Topic 5)
You are deploying an application to Google Cloud. The application is part of a system. The application in Google Cloud must communicate over a private network with applications in a non-Google Cloud environment. The expected average throughput is 200 kbps. The business requires:
• 99.99% system availability
• cost optimization
You need to design the connectivity between the locations to meet the business requirements. What should you provision?

A. A Classic Cloud VPN gateway connected with one tunnel to an on-premises VPN gateway.
B. A Classic Cloud VPN gateway connected with two tunnels to an on-premises VPN gateway.
C. An HA Cloud VPN gateway connected with two tunnels to an on-premises VPN gateway.
D. Two HA Cloud VPN gateways connected to two on-premises VPN gateway
E. Configure each HA CloudVPN gateway to have two tunnels, each connected to different on-premises VPN gateways.

Answer: C

Explanation:
https://cloud.google.com/network-connectivity/docs/vpn/concepts/topologies#configurations_that_support_9999_availability

NEW QUESTION 74
- (Topic 5)
Your company uses Google Kubernetes Engine (GKE) as a platform for all workloads. Your company has a single large GKE cluster that contains batch, stateful, and stateless workloads. The GKE cluster is configured with a single node pool with 200 nodes. Your company needs to reduce the cost of this cluster but does not want to compromise availability. What should you do?

A. Create a second GKE cluster for the batch workloads onl
B. Allocate the 200 original nodes across both clusters.
C. Configure a HorizontalPodAutoscaler for all stateless workloads and for all compatible stateful workload
D. Configure the cluster to use node auto scaling.
E. Configure CPU and memory limits on the namespaces in the cluste
F. Configure all Pods to have a CPU and memory limits.
G. Change the node pool to use spot VMs.

Answer: C

Explanation:
One way to reduce the cost of a Google Kubernetes Engine (GKE) cluster without compromising availability is to use horizontal pod autoscalers (HPA) and node auto scaling. HPA allows you to automatically scale the number of Pods in a deployment based on the resource usage of the Pods. By configuring HPA for stateless workloads and for compatible stateful workloads, you can ensure that the number of Pods is automatically adjusted based on the actual resource usage, which can help to reduce costs. Node auto scaling allows you to automatically add or remove nodes from the node pool based on the resource usage of the cluster. By configuring node auto scaling, you can ensure that the cluster has the minimum number of nodes needed to meet the resource requirements of the

workloads, which can also help to reduce costs.

NEW QUESTION 75
- (Topic 5)
Your company and one of its partners each nave a Google Cloud protect in separate organizations. Your company s protect (prj-a) runs in Virtual Private Cloud (vpc-a). The partner's project (prj-b) runs in vpc-b. There are two instances running on vpc-a and one instance running on vpc-b Subnets denned in both VPCs are not overlapping. You need to ensure that all instances communicate with each other via internal IPs minimizing latency and maximizing throughput. What should you do?

A. Set up a network peering between vpc-a and vpc-b
B. Set up a VPN between vpc-a and vpc-b using Cloud VPN
C. Configure IAP TCP forwarding on the instance in vpc b and then launch the following gcloud command from one of the instance in vpc-gcloud:
* 1. Create an additional instance in vpc-a* 2. Create an additional instance n vpc-b* 3. Instal OpenVPN in newly created instances* 4. Configure a VPN tunnel between vpc-a and vpc-b with the help of OpenVPN

Answer: C

NEW QUESTION 79
- (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 transferarchived data to Cloud Storag
C. Establish a connection with Google using a Dedicated Interconnect orDirect 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 84
- (Topic 5)
Your company is running a stateless application on a Compute Engine instance. The application is used
heavily during regular business hours and lightly outside of business hours. Users are reporting that the application is slow during peak hours. You need to optimize the application's performance. What should you do?

A. Create a snapshot of the existing dis
B. Create an instance template from the snapshot.Create anautoscaled managed instance group from the instance template.
C. Create a snapshot of the existing dis
D. Create a custom image from the snapsho
E. Create an autoscaled managed instance group from the custom image.
F. Create a custom image from the existing dis
G. Create an instance template from the custom imag
H. Create an autoscaled managed instance group from the instance template.
I. Create an instance template from the existing dis
J. Create a custom image from the instance template.Create an autoscaled managed instance group from the custom image.

Answer: B

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

NEW QUESTION 88
- (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 93
- (Topic 5)
Your company has a stateless web API that performs scientific calculations. The web API runs on a single Google Kubernetes Engine (GKE) cluster. The cluster is currently deployed in us-central1. Your company has expanded to offer your API to customers in Asia. You want to reduce the latency for the users in Asia. What should you do?

A. Use a global HTTP(s) load balancer with Cloud CDN enabled
B. Create a second GKE cluster in asia-southeast1, and expose both API's using a Service oftype Load Balance
C. Add the public Ips to the Cloud DNS zone
D. Increase the memory and CPU allocated to the application in the cluster
E. Create a second GKE cluster in asia-southeast1, and use kubemci to create a global HTTP(s) load balancer

Answer: D

Explanation:
 https://cloud.google.com/kubernetes-engine/docs/concepts/multi-cluster- ingress#how_works
https://github.com/GoogleCloudPlatform/k8s-multicluster-ingress https://cloud.google.com/blog/products/gcp/how-to-deploy-geographically-distributed-services-on-kubernetes-engine-with-kubemci

NEW QUESTION 95
- (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 99
- (Topic 5)
You have been asked to select the storage system for the click-data of your company's large portfolio of websites. This data is streamed in from a custom website analytics package at a typical rate of 6,000 clicks per minute, with bursts of up to 8,500 clicks per second. It must been stored for future analysis by your data science and user experience teams. Which storage infrastructure should you choose?

A. Google Cloud SQL
B. Google Cloud Bigtable
C. Google Cloud Storage
D. Google cloud Datastore

Answer: C

Explanation:
https://cloud.google.com/bigquery/docs/loading-data-cloud-storage

NEW QUESTION 102
- (Topic 5)
For this question, refer to the TerramEarth case study. You are building a microservice- based application for TerramEarth. The application is based on Docker containers. You want to follow Google-recommended practices to build the application continuously and store the build artifacts. What should you do?
A.
* 1. Configure a trigger in Cloud Build for new source changes.
* 2. Invoke Cloud Build to build one container image, and tag the image with the label 'latest.'
* 3. Push the image to the Artifact Registry.
B.
* 1. Configure a trigger in Cloud Build for new source changes.
* 2. Invoke Cloud Build to build container images for each microservice, and tag them using the code commit hash.
* 3. Push the images to the Artifact Registry.
C.
* 1 Create a Scheduler job to check the repo every minute.
* 2. For any new change, invoke Cloud Build to build container images for the microservices.
* 3. Tag the images using the current timestamp, and push them to the Artifact Registry.
D.
* 1. Configure a trigger in Cloud Build for new source changes.
* 2. The trigger invokes build jobs and build container images for the microservices.
* 3. Tag the images with a version number, and push them to Cloud Storage.

A.

Answer: C

NEW QUESTION 107
- (Topic 5)
You need to develop procedures to test a disaster plan for a mission-critical application. You want to use
Google-recommended practices and native capabilities within GCP. What should you do?

A. Use Deployment Manager to automate service provisionin
B. Use Activity Logs to monitor and debug your tests.
C. Use Deployment Manager to automate provisionin
D. Use Stackdriver to monitor and debug your tests.
E. Use gcloud scripts to automate service provisionin
F. Use Activity Logs monitor and debug your tests.
G. Use automated scripts to automate service provisionin
H. Use Activity Logs monitor and debug your tests.

Answer: B

Explanation:
https://cloud.google.com/solutions/dr-scenarios-planning-guide

NEW QUESTION 112
- (Topic 5)
You want to automate the creation of a managed instance group and a startup script to install the OS package dependencies. You want to minimize the startup
time for VMs in the instance group.
What should you do?

A. Use Terraform to create the managed instance group and a startup script to install the OS packagedependencies.
B. Create a custom VM image with all OS package dependencie
C. Use Deployment Manager to create the managed instance group with the VM image.
D. Use Puppet to create the managed instance group and install the OS package dependencies.
E. Use Deployment Manager to create the managed instance group and Ansible to install the OS package dependencies.

Answer: B

Explanation:
"Custom images are more deterministic and start more quickly than instances with startup scripts. However, startup scripts are more flexible and let you update the
apps and settings in your instances more easily." https://cloud.google.com/compute/docs/instance- templates/create-instance-
templates#using_custom_or_public_images_in_your_instance_templates

NEW QUESTION 114
- (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 118
- (Topic 5)
You need to deploy an application on Google Cloud that must run on a Debian Linux environment. The application requires extensive configuration in order to
operate correctly. You want to ensure that you can install Debian distribution updates with minimal manual intervention whenever they become available. What
should you do?

A. Create a Compute Engine instance template using the most recent Debian imag
B. Create an instance from this template, and install and configure the application as part of the startup scrip
C. Repeat this process whenever a new Google-managed Debian image becomes available.
D. Create a Debian-based Compute Engine instance, install and configure the application, and use OS patch management to install available updates.
E. Create an instance with the latest available Debian imag
F. Connect to the instance via SSH, and install and configure the application on the instanc
G. Repeat this process whenever a new Google-managed Debian image becomes available.
H. Create a Docker container with Debian as the base imag
I. Install and configure the application as part of the Docker image creation proces
J. Host the container on Google Kubernetes Engine and restart the container whenever a new update is available.

Answer: B

Explanation:
Reference: https://cloud.google.com/compute/docs/os-patch-management

NEW QUESTION 119
- (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 andconfigure 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 122
- (Topic 5)
Your company has an application deployed on Anthos clusters (formerly Anthos GKE) that is running multiple microservices. The cluster has both Anthos Service Mesh and Anthos Config Management configured. End users inform you that the application is responding very slowly. You want to identify the microservice that is causing the delay. What should you do?

A. Use the Service Mesh visualization in the Cloud Console to inspect the telemetry between the microservices.
B. Use Anthos Config Management to create a ClusterSelector selecting the relevant cluste
C. On the Google Cloud Console page for Google Kubernetes Engine, view the Workloads and filter on the cluste
D. Inspect the configurations of the filtered workloads.
E. Use Anthos Config Management to create a namespaceSelector selecting the relevantcluster namespac
F. On the Google Cloud Console page for Google Kubernetes Engine, visit the workloads and filter on the namespac
G. Inspect the configurations of the filtered workloads.
H. Reinstall istio using the default istio profile in order to collect request latenc
I. Evaluate the telemetry between the microservices in the Cloud Console.

Answer: A

Explanation:
 The Anthos Service Mesh pages in the Google Cloud Console provide both summary and in-depth metrics, charts, and graphs that enable you to observe service behavior. You can monitor the overall health of your services, or drill down on a specific service to set a service level objective (SLO) or troubleshoot an issue. https://cloud.google.com/service-mesh/docs/observability/explore-dashboard https://cloud.google.com/anthos/service-mesh

NEW QUESTION 127
- (Topic 5)
You write a Python script to connect to Google BigQuery from a Google Compute Engine virtual machine. The script is printing errors that it cannot connect to BigQuery. What should you do to fix the script?

A. Install the latest BigQuery API client library for Python
B. Run your script on a new virtual machine with the BigQuery access scope enabled
C. Create a new service account with BigQuery access and execute your script with that user
D. Install the bq component for gccloud with the command gcloud components install bq.

Answer: B

Explanation:
The error is most like caused by the access scope issue. When create new instance, you have the default Compute engine default service account but most serves access including BigQuery is not enable. Create an instance Most access are not enabled by default You have default service account but don't have the permission (scope) you can stop the instance, edit, change scope and restart it to enable the scope access. Of course, if you Run your script on a new virtual machine with the BigQuery access scope enabled, it also works
https://cloud.google.com/compute/docs/access/service-accounts

NEW QUESTION 128
- (Topic 5)
You are managing several internal applications that are deployed on Compute Engine. Business users inform you that an application has become very slow over the past few days. You want to find the underlying cause in order to solve the problem. What should you do first?

A. Inspect the logs and metrics from the instances in Cloud Logging and Cloud Monitoring.
B. Restore a backup of the application database from a time before the application became slow.
C. Deploy the applications on a managed instance group with autoscaling enable
D. Add aload balancer in front of the managed instance group, and have the users connect to the IP of the load balancer.
E. Change the Compute Engine Instances behind the application to a machine type with more CPU and memory.

Answer: A

Explanation:
 When an application becomes slow, the first step you should take is to gather information about the underlying cause of the problem. One way to do this is by inspecting the logs and metrics from the instances where the application is deployed. Google Cloud Platform (GCP) provides tools such as Cloud Logging and Cloud Monitoring that can help you to collect and analyze this information. By reviewing the logs and metrics from the instances, you may be able to identify issues such as resource shortages (e.g. CPU, memory, or disk), network problems, or application errors that are causing the performance issues. Once you have identified the underlying cause of the problem, you can take steps to resolve it.

NEW QUESTION 129
- (Topic 5)
You company has a Kubernetes application that pulls messages from Pub/Sub and stores them in Firestore. Because the application is simple, it was deployed as a single pod. The infrastructure team has analyzed Pub/Sub metrics and discovered that the application cannot process the messages in real time. Most of them wait for minutes before being processed. You need to scale the elaboration process that is I/O-intensive. What should you do?

A. Configure a Kubernetes autoscaling based on the subscription/push_request metric.
B. Use the –enable- autoscaling flag when you create the Kubernetes cluster
C. Configure a Kubernetes autoscaling based on the subscription/num_undelivered message metric.
D. Use kubectl autoscale deployment APP_NAME –max 6 –min 2 –cpu- percent 50 to configure Kubernetes autoscaling deployment

Answer: A

Explanation:
 https://cloud.google.com/kubernetes-engine/docs/concepts/custom-and-external-metrics#external_metrics

NEW QUESTION 134
- (Topic 5)
Your company has an application running on Google Cloud that is collecting data from thousands of physical devices that are globally distributed. Data is publish to Pub/Sub and streamed in real time into an SSO Cloud Bigtable cluster via a Dataflow pipeline. The operations team informs you that your Cloud Bigtable cluster has a hot-spot and queries are taking longer man expected You need to resolve the problem and prevent it from happening in the future What should you do?

A. Advise your clients to use HBase APIs instead of NodeJS APIs.
B. Review your RowKey strategy and ensure that keys are evenly spread across the alphabet.
C. Delete records older than 30 days.
D. Double the number of nodes you currently have.

Answer: B

NEW QUESTION 136
- (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 138
- (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 139
- (Topic 5)
Your company has decided to build a backup replica of their on-premises user authentication PostgreSQL database on Google Cloud Platform. The database is 4 TB, and large updates are frequent. Replication requires private address space communication. Which networking approach should you use?

A. Google Cloud Dedicated Interconnect
B. Google Cloud VPN connected to the data center network
C. A NAT and TLS translation gateway installed on-premises
D. A Google Compute Engine instance with a VPN server installed connected to the data center network

Answer: A

Explanation:

https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations
Google Cloud Dedicated Interconnect provides direct physical connections and RFC 1918 communication between your on-premises network and Google's network. Dedicated Interconnect enables you to transfer large amounts of data between networks, which can be more cost effective than purchasing additional bandwidth over the public Internet or using VPN tunnels.
Benefits:
? Traffic between your on-premises network and your VPC network doesn't traverse the public Internet. Traffic traverses a dedicated connection with fewer hops, meaning there are less points of failure where traffic might get dropped or disrupted.
? Your VPC network's internal (RFC 1918) IP addresses are directly accessible from your on-premises network. You don't need to use a NAT device or VPN tunnel to reach internal IP addresses. Currently, you can only reach internal IP addresses over a dedicated connection. To reach Google external IP addresses, you must use a separate connection.
? You can scale your connection to Google based on your needs. Connection capacity is delivered over one or more 10 Gbps Ethernet connections, with a maximum of eight connections (80 Gbps total per interconnect).
? The cost of egress traffic from your VPC network to your on-premises network is reduced. A dedicated connection is generally the least expensive method if you have a high-volume of traffic to and from Google's network.
References: https://cloud.google.com/interconnect/docs/details/dedicated

NEW QUESTION 144
- (Topic 5)
Your architecture calls for the centralized collection of all admin activity and VM system logs within your
project.
How should you collect these logs from both VMs and services?

A. All admin and VM system logs are automatically collected by Stackdriver.
B. Stackdriver automatically collects admin activity logs for most service
C. The Stackdriver Logging agentmust be installed on each instance to collect system logs.
D. Launch a custom syslogd compute instance and configure your GCP project and VMs to forward all logs to it.
E. Install the Stackdriver Logging agent on a single compute instance and let it collect all audit and access logs for your environment.

Answer: B

Explanation:
https://cloud.google.com/logging/docs/agent/default-logs

NEW QUESTION 145
- (Topic 5)
You are designing a large distributed application with 30 microservices. Each of your distributed microservices needs to connect to a database back-end. You want to store the credentials securely. Where should you store the credentials?

A. In the source code
B. In an environment variable
C. In a secret management system
D. In a config file that has restricted access through ACLs

Answer: C

Explanation:
https://cloud.google.com/docs/authentication/production#providing_credentials_to_your_ap plication

NEW QUESTION 147
- (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 investigat
E. 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 152
- (Topic 5)
Your company has sensitive data in Cloud Storage buckets. Data analysts have Identity
Access Management (IAM) permissions to read the buckets. You want to prevent data analysts from retrieving the data in the buckets from outside the office network. What should you do?

A. * 1. Create a VPC Service Controls perimeter that includes the projects with the buckets.* 2. Create an access level with the CIDR of the office network.
B. * 1. Create a firewall rule for all instances in the Virtual Private Cloud (VPC) network for source range.* 2. Use the Classless Inter-domain Routing (CIDR) of the office network.
C. * 1. Create a Cloud Function to remove IAM permissions from the buckets, and another Cloud Function to add IAM permissions to the buckets.* 2. Schedule the Cloud Functions with Cloud Scheduler to add permissions at the start of business and remove permissions at the end of business.
D. * 1. Create a Cloud VPN to the office network.* 2. Configure Private Google Access for on-premises hosts.

Answer: A

Explanation:
For all Google Cloud services secured with VPC Service Controls, you can ensure that: Resources within a perimeter are accessed only from clients within authorized VPC networks using Private Google Access with either Google Cloud or on-premises. https://cloud.google.com/vpc-service-controls/docs/overview https://cloud.google.com/vpc-service-controls/docs/overview. You create a service control
across your VPC and any cloud bucket or any project resource to restrict access. Anything outside of it can't access the resources within service control perimeter

NEW QUESTION 154
- (Topic 5)
Your organization has a 3-tier web application deployed in the same network on Google Cloud Platform. Each tier (web, API, and database) scales independently of the others Network traffic should flow through the web to the API tier and then on to the database tier. Traffic should not flow between the web and the database tier. How should you configure the network?

A. Add each tier to a different subnetwork.
B. Set up software based firewalls on individual VMs.
C. Add tags to each tier and set up routes to allow the desired traffic flow.
D. Add tags to each tier and set up firewall rules to allow the desired traffic flow.

Answer: D

Explanation:
 https://aws.amazon.com/blogs/aws/building-three-tier-architectures-with- security-groups/
Google Cloud Platform(GCP) enforces firewall rules through rules and tags. GCP rules and tags can be defined once and used across all regions.
References: https://cloud.google.com/docs/compare/openstack/ https://aws.amazon.com/it/blogs/aws/building-three-tier-architectures-with-security-groups/

NEW QUESTION 159
- (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 costmanagement.
C. Utilize free tier and sustained use discount
D. Provision a staff position for 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 160
- (Topic 6)
For this question, refer to the Dress4Win case study. Which of the compute services should be migrated as –is and would still be an optimized architecture for performance in the cloud?

A. Web applications deployed using App Engine standard environment
B. RabbitMQ deployed using an unmanaged instance group
C. Hadoop/Spark deployed using Cloud Dataproc Regional in High Availability mode
D. Jenkins, monitoring, bastion hosts, security scanners services deployed on custom machine types

Answer: C

NEW QUESTION 162
- (Topic 7)
You have broken down a legacy monolithic application into a few containerized RESTful microservices. You want to run those microservices on Cloud Run. You also want to make sure the services are highly available with low latency to your customers. What should you do?

A. Deploy Cloud Run services to multiple availability zone
B. Create Cloud Endpoints that point to the service
C. Create a global HTIP(S) Load Balancing instance and attach the Cloud Endpoints to its backend.
D. Deploy Cloud Run services to multiple regions Create serverless network endpoint groups pointing to the service
E. Add the serverless NE Gs to a backend service that is used by a global HTIP(S) Load Balancing instance.
F. Cloud Run services to multiple region
G. In Cloud DNS, create a latency-based DNS name that points to the services.
H. Deploy Cloud Run services to multiple availability zone
I. Create a TCP/IP global load balance
J. Add the Cloud Run Endpoints to its backend service.

Answer: B

Explanation:
https://cloud.google.com/run/docs/multiple-regions

NEW QUESTION 163
- (Topic 7)
For this question, refer to the TerramEarth case study. To be compliant with European GDPR regulation, TerramEarth is required to delete data generated from its European customers after a period of 36 months when it contains personal data. In the new architecture, this data will be stored in both Cloud Storage and BigQuery. What should you do?

A. Create a BigQuery table for the European data, and set the table retention period to 36 month
B. For Cloud Storage, use gsutil to enable lifecycle management using a DELETE action with an Age condition of 36 months.
C. Create a BigQuery table for the European data, and set the table retention period to 36 month
D. For Cloud Storage, use gsutil to create a SetStorageClass to NONE action when with an Age condition of 36 months.
E. Create a BigQuery time-partitioned table for the European data, and set the partition expiration period to 36 month
F. For Cloud Storage, use gsutil to enable lifecycle management using a DELETE action with an Age condition of 36 months.
G. Create a BigQuery time-partitioned table for the European data, and set the partition period to 36 month
H. For Cloud Storage, use gsutil to create a SetStorageClass to NONE action with an Age condition of 36 months.

Answer: C

Explanation:
 https://cloud.google.com/bigquery/docs/managing-partitioned- tables#partition-expiration
https://cloud.google.com/storage/docs/lifecycle

NEW QUESTION 166
- (Topic 7)
For this question, refer to the TerramEarth case study. You are asked to design a new architecture for the
ingestion of the data of the 200,000 vehicles that are connected to a cellular network. You want to follow
Google-recommended practices.
Considering the technical requirements, which components should you use for the ingestion of the data?

A. Google Kubernetes Engine with an SSL Ingress
B. Cloud IoT Core with public/private key pairs
C. Compute Engine with project-wide SSH keys
D. Compute Engine with specific SSH keys

Answer: A

Explanation:
 https://cloud.google.com/solutions/iot-overview https://cloud.google.com/iot/quotas

NEW QUESTION 171
- (Topic 8)
For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to migrate from their current analytics and statistics reporting model to one
that meets their technical requirements on Google Cloud Platform.
Which two steps should be part of their migration plan? (Choose two.)

A. Evaluate the impact of migrating their current batch ETL code to Cloud Dataflow.
B. Write a schema migration plan to denormalize data for better performance in BigQuery.
C. Draw an architecture diagram that shows how to move from a single MySQL database to a MySQL cluster.
D. Load 10 TB of analytics data from a previous game into a Cloud SQL instance, and run test queries against the full dataset to confirm that they complete
successfully.
E. Integrate Cloud Armor to defend against possible SQL injection attacks in analytics files uploaded to Cloud Storage.

Answer: AB

Explanation:
 https://cloud.google.com/bigquery/docs/loading-data#loading_denormalized_nested_and_repeated_data

NEW QUESTION 172
- (Topic 8)
You need to implement a network ingress for a new game that meets the defined business and technical
requirements. Mountkirk Games wants each regional game instance to be located in multiple Google Cloud
regions. What should you do?

A. Configure a global load balancer connected to a managed instance group running Compute Engineinstances.
B. Configure kubemci with a global load balancer and Google Kubernetes Engine.
C. Configure a global load balancer with Google Kubernetes Engine.
D. Configure Ingress for Anthos with a global load balancer and Google Kubernetes Engine.

Answer: A

NEW QUESTION 175
- (Topic 8)
Mountkirk Games wants to limit the physical location of resources to their operating Google Cloud regions.
What should you do?

A. Configure an organizational policy which constrains where resources can be deployed.
B. Configure IAM conditions to limit what resources can be configured.
C. Configure the quotas for resources in the regions not being used to 0.
D. Configure a custom alert in Cloud Monitoring so you can disable resources as they are created in otherregions.

Answer: A

NEW QUESTION 177
- (Topic 9)
For this question, refer to the Helicopter Racing League (HRL) case study. HRL wants better prediction
accuracy from their ML prediction models. They want you to use Google's AI Platform so HRL can understand and interpret the predictions. What should you do?

A. Use Explainable AI.
B. Use Vision AI.
C. Use Google Cloud's operations suite.
D. Use Jupyter Notebooks.

Answer: A

Explanation:
 Reference: https://cloud.google.com/ai-platform/prediction/docs/ai-explanations/preparing- metadata

NEW QUESTION 180
- (Topic 9)
For this question, refer to the Helicopter Racing League (HRL) case study. A recent finance audit of cloud
infrastructure noted an exceptionally high number of Compute Engine instances are allocated to do video
encoding and transcoding. You suspect that these Virtual Machines are zombie machines that were not deleted after their workloads completed. You need to quickly get a list of which VM instances are idle. What should you do?

A. Log into each Compute Engine instance and collect disk, CPU, memory, and network usage statistics foranalysis.
B. Use the gcloud compute instances list to list the virtual machine instances that have the idle: true label set.
C. Use the gcloud recommender command to list the idle virtual machine instances.
D. From the Google Console, identify which Compute Engine instances in the managed instance groups areno longer responding to health check probes.

Answer: C

Explanation:
 Reference: https://cloud.google.com/compute/docs/instances/viewing-and-applying-idle- vm-recommendations

NEW QUESTION 182
- (Topic 10)
For this question, refer to the EHR Healthcare case study. In the past, configuration errors put public IP addresses on backend servers that should not have been accessible from the Internet. You need to ensure that no one can put external IP addresses on backend Compute Engine instances and that external IP addresses can only be configured on frontend Compute Engine instances. What should you do?

A. Create an Organizational Policy with a constraint to allow external IP addresses only on the frontend Compute Engine instances.
B. Revoke the compute.networkAdmin role from all users in the project with front end instances.
C. Create an Identity and Access Management (IAM) policy that maps the IT staff to the compute.networkAdmin role for the organization.
D. Create a custom Identity and Access Management (IAM) role named GCE_FRONTEND with the compute.addresses.create permission.

Answer: A

Explanation:
 https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#disableexternalip

NEW QUESTION 187
- (Topic 10)
For this question, refer to the EHR Healthcare case study. You are responsible for ensuring that EHR's use of Google Cloud will pass an upcoming privacy compliance audit. What should you do? (Choose two.)

A. Verify EHR's product usage against the list of compliant products on the Google Cloud compliance page.
B. Advise EHR to execute a Business Associate Agreement (BAA) with Google Cloud.
C. Use Firebase Authentication for EHR's user facing applications.
D. Implement Prometheus to detect and prevent security breaches on EHR's web-based applications.
E. Use GKE private clusters for all Kubernetes workloads.

Answer: AB

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

NEW QUESTION 188
- (Topic 10)
For this question, refer to the EHR Healthcare case study. You are responsible for designing the Google Cloud network architecture for Google Kubernetes
Engine. You want to follow Google best practices. Considering the EHR Healthcare business and technical requirements, what should you do to reduce the attack surface?

A. Use a private cluster with a private endpoint with master authorized networks configured.
B. Use a public cluster with firewall rules and Virtual Private Cloud (VPC) routes.
C. Use a private cluster with a public endpoint with master authorized networks configured.
D. Use a public cluster with master authorized networks enabled and firewall rules.

Answer: A

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/private-cluster- concept#overview

NEW QUESTION 189
- (Topic 10)
For this question, refer to the EHR Healthcare case study. You need to define the technical architecture for hybrid connectivity between EHR's on-premises systems and Google Cloud. You want to follow Google's recommended practices for production-level applications. Considering the EHR Healthcare business and technical requirements, what should you do?

A. Configure two Partner Interconnect connections in one metro (City), and make sure the Interconnect connections are placed in different metro zones.
B. Configure two VPN connections from on-premises to Google Cloud, and make sure the VPN devices on-premises are in separate racks.
C. Configure Direct Peering between EHR Healthcare and Google Cloud, and make sure you are peering at least two Google locations.
D. Configure two Dedicated Interconnect connections in one metro (City) and two connections in another metro, and make sure the Interconnect connections are placed in different metro zones.

Answer: D

Explanation:
based on the requirement of secure and high-performance connection between on-premises systems to Google Cloud
https://cloud.google.com/network-connectivity/docs/interconnect/tutorials/partner-creating- 9999-availability

NEW QUESTION 191
- (Topic 10)
For this question, refer to the EHR Healthcare case study. You need to define the technical architecture for securely deploying workloads to Google Cloud. You also need to ensure that only verified containers are deployed using Google Cloud services. What should you do? (Choose two.)

A. Enable Binary Authorization on GKE, and sign containers as part of a CI/CD pipeline.
B. Configure Jenkins to utilize Kritis to cryptographically sign a container as part of a CI/CD pipeline.
C. Configure Container Registry to only allow trusted service accounts to create and deploy containers from the registry.
D. Configure Container Registry to use vulnerability scanning to confirm that there are no vulnerabilities before deploying the workload.

Answer: A

Explanation:
Binary Authorization to ensure only verified containers are deployed To ensure deployment are secure and and consistent, automatically scan images for vulnerabilities with container analysis (https://cloud.google.com/docs/ci-cd/overview?hl=en&skip_cache=true)

NEW QUESTION 193
......

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

    All transactions are protected by VeriSign!

100% Pass Your Professional-Cloud-Architect Exam with Our Prep Materials Via below:

https://www.certleader.com/Professional-Cloud-Architect-dumps.html