Get Google Professional-Cloud-Architect Dumps Questions [2023 To Gain Brilliant Result [Q112-Q129

Get Google Professional-Cloud-Architect Dumps Questions [2023 To Gain Brilliant Result Professional-Cloud-Architect dumps - Dumpleader - 100% Passing Guarantee

Google Professional-Cloud-Architect certification exam is designed to test an individual's knowledge and skills in cloud architecture and Google Cloud Platform (GCP). It is intended for professionals who work with GCP and have experience designing, developing, and managing cloud solutions.

To prepare for the GCP exam, candidates can take advantage of a range of resources available online. Google offers a range of training programs, including online courses, tutorials, and hands-on labs, to help candidates build their knowledge and skills in GCP. Additionally, third-party training providers offer courses and study materials that can help candidates prepare for the exam.

Section #5. Managing implementation

This section has two major topics and assesses the candidates' skills in application development, system migration, data management, API development/usage best practices, and testing frameworks. Interacting with Google Cloud programmatically is the key focus of the second sub-topic. The main considerations here are Cloud Emulators, Google Cloud Shell, and Cloud SDK.

Q112. You are implementing Firestore for Mountkirk Games. Mountkirk Games wants to give a new game programmatic access to a legacy game’s Firestore database. Access should be as restricted as possible. What should you do?
* Create a service account (SA) in the legacy game’s Google Cloud project, add a second SA in the new game’s IAM page, and then give the Organization Admin role to both SAs.
* Create a service account (SA) in the legacy game’s Google Cloud project, give the SA the Organization Admin role, and then give it the Firebase Admin role in both projects.
* Create a service account (SA) in the legacy game’s Google Cloud project, add this SA in the new game’s IAM page, and then give it the Firebase Admin role in both projects.
* Create a service account (SA) in the legacy game’s Google Cloud project, give it the Firebase Admin role, and then migrate the new game to the legacy game’s project.

Q113. Your company wants to track whether someone is present in a meeting room reserved for a scheduled meeting. There are 1000 meeting rooms across 5 offices on 3 continents. Each room is equipped with a motion sensor that reports its status every second. The data from the motion detector includes only a sensor ID and several different discrete items of information. Analysts will use this data, together with information about account owners and office locations. Which database type should you use?
* Flat file
* NoSQL
* Relational
* Blobstore

Q114. You want to enable your running Google Kubernetes Engine cluster to scale as demand for your application changes.

What should you do?
* Add additional nodes to your Kubernetes Engine cluster using the following command:

gcloud container clusters resize

CLUSTER_Name – -size 10
* Add a tag to the instances in the cluster with the following command:

gcloud compute instances add-tags

INSTANCE – -tags enable-

autoscaling max-nodes-10
* Update the existing Kubernetes Engine cluster with the following command:

gcloud alpha container clusters

update mycluster – -enable-

autoscaling – -min-nodes=1 – -max-nodes=10
* Create a new Kubernetes Engine cluster with the following command:

gcloud alpha container clusters

create mycluster – -enable-

autoscaling – -min-nodes=1 – -max-nodes=10

and redeploy your application

Q115. 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?
* Google Kubernetes Engine, Jenkins, and Helm
* Google Kubernetes Engine and Cloud Load Balancing
* Google Kubernetes Engine and Cloud Deployment Manager
* Google Kubernetes Engine, Jenkins, and Cloud Load Balancing

Explanation/Reference:

Explanation:

Jenkins is an open-source automation server that lets you flexibly orchestrate your build, test, and

deployment pipelines. Kubernetes Engine is a hosted version of Kubernetes, a powerful cluster manager

and orchestration system for containers.

When you need to set up a continuous delivery (CD) pipeline, deploying Jenkins on Kubernetes Engine

provides important benefits over a standard VM-based deployment

Incorrect Answers:

A: Helm is a tool for managing Kubernetes charts. Charts are packages of pre-configured Kubernetes

resources.

Use Helm to:

Find and use popular software packaged as Kubernetes charts

▪

Share your own applications as Kubernetes charts

▪

Create reproducible builds of your Kubernetes applications

▪

Intelligently manage your Kubernetes manifest files

▪

Manage releases of Helm packages

▪

References: https://cloud.google.com/solutions/jenkins-on-kubernetes-engine

Q116. Your company has multiple on-premises systems that serve as sources for reporting. The data has not been maintained well and has become degraded over time. You want to use Google-recommended practices to detect anomalies in your company dat a. What should you do?
* Upload your files into Cloud Storage. Use Cloud Datalab to explore and clean your data.
* Upload your files into Cloud Storage. Use Cloud Dataprep to explore and clean your data.
* Connect Cloud Datalab to your on-premises systems. Use Cloud Datalab to explore and clean your data.
* Connect Cloud Dataprep to your on-premises systems. Use Cloud Dataprep to explore and clean your data.

Reference:

https://cloud.google.com/dataprep/

Q117. You want to enable your running Google Container Engine cluster to scale as demand for your application changes.

What should you do?
* Add additional nodes to your Container Engine cluster using the following command:

gcloud container clusters resize CLUSTER_NAME –size 10
* Add a tag to the instances in the cluster with the following command:

gcloud compute instances add-tags INSTANCE –tags enable –autoscaling max-nodes-10
* Update the existing Container Engine cluster with the following command:

gcloud alpha container clusters update mycluster –enable-autoscaling –min-nodes=1 –max-nodes=10
* Create a new Container Engine cluster with the following command:

gcloud alpha container clusters create mycluster –enable-autocaling –min-nodes=1 –max-nodes=10 and redeploy your application.

Explanation:

Cluster autoscaling

–enable-autoscaling

Enables autoscaling for a node pool.

Enables autoscaling in the node pool specified by –node-pool or the default node pool if –node-pool is not provided.

Where:

–max-nodes=MAX_NODES

Maximum number of nodes in the node pool.

Maximum number of nodes to which the node pool specified by –node-pool (or default node pool if unspecified) can scale.

Incorrect Answers:

C, D: Warning: Do not use Alpha Clusters or alpha features for production workloads.

Note: You can experiment with Kubernetes alpha features by creating an alpha cluster. Alpha clusters are short-lived clusters that run stable Kubernetes releases with all Kubernetes APIs and features enabled. Alpha clusters are designed for advanced users and early adopters to experiment with workloads that take advantage of new features before those features are production-ready. You can use Alpha clusters just like normal Kubernetes Engine clusters.

References: https://cloud.google.com/sdk/gcloud/reference/container/clusters/create Reference:

https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler

Q118. 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?
* Direct them to download and install the Google StackDriver logging agent.
* Send them a list of online resources about logging best practices.
* Help them define their requirements and assess viable logging tools.
* Help them upgrade their current tool to take advantage of any new features.

Explanation

The Stackdriver Logging agent streams logs from your VM instances and from selected third party software packages to Stackdriver Logging. Using the agent is optional but we recommend it. The agent runs under both Linux and Microsoft Windows.

Note: 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/docs/agent/installation

Q119. A development team at your company has created a dockerized HTTPS web application. You need to deploy the application on Google Kubernetes Engine (GKE) and make sure that the application scales automatically.

How should you deploy to GKE?
* Use the Horizontal Pod Autoscaler and enable cluster autoscaling. Use an Ingress resource to loadbalance the HTTPS traffic.
* Use the Horizontal Pod Autoscaler and enable cluster autoscaling on the Kubernetes cluster. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.
* Enable autoscaling on the Compute Engine instance group. Use an Ingress resource to load balance the HTTPS traffic.
* Enable autoscaling on the Compute Engine instance group. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.
Explanation

Reference https://cloud.google.com/kubernetes-engine/docs/how-to/cluster-autoscaler

Q120. Case Study: 4 – Dress4Win case study

Company Overview

Dress4win is a web-based company that helps their users organize and manage their personal wardrobe using a website and mobile application. The company also cultivates an active social network that connects their users with designers and retailers. They monetize their services through advertising, e-commerce, referrals, and a freemium app model.

Company Background

Dress4win’s application has grown from a few servers in the founder’s garage to several hundred servers and appliances in a colocated data center. However, the capacity of their infrastructure is now insufficient for the application’s rapid growth. Because of this growth and the company’s desire to innovate faster, Dress4win is committing to a full migration to a public cloud.

Solution Concept

For the first phase of their migration to the cloud, Dress4win is considering moving their development and test environments. They are also considering building a disaster recovery site, because their current infrastructure is at a single location. They are not sure which components of their architecture they can migrate as is and which components they need to change before migrating them.

Existing Technical Environment

The Dress4win application is served out of a single data center location.

Databases:

MySQL – user data, inventory, static data

* Redis – metadata, social graph, caching

* Application servers:

Tomcat – Java micro-services

* Nginx – static content

* Apache Beam – Batch processing

* Storage appliances:

iSCSI for VM hosts

* Fiber channel SAN – MySQL databases

* NAS – image storage, logs, backups

* Apache Hadoop/Spark servers:

Data analysis

* Real-time trending calculations

* MQ servers:

Messaging

* Social notifications

* Events

* Miscellaneous servers:

Jenkins, monitoring, bastion hosts, security scanners

* Business Requirements

* Build a reliable and reproducible environment with scaled parity of production. Improve security by defining and adhering to a set of security and Identity and Access Management (IAM) best practices for cloud.

Improve business agility and speed of innovation through rapid provisioning of new resources.

Analyze and optimize architecture for performance in the cloud. Migrate fully to the cloud if all other requirements are met.

Technical Requirements

Evaluate and choose an automation framework for provisioning resources in cloud. Support failover of the production environment to cloud during an emergency. Identify production services that can migrate to cloud to save capacity.

Use managed services whenever possible.

Encrypt data on the wire and at rest.

Support multiple VPN connections between the production data center and cloud environment.

CEO Statement

Our investors are concerned about our ability to scale and contain costs with our current infrastructure. They are also concerned that a new competitor could use a public cloud platform to offset their up-front investment and freeing them to focus on developing better features.

CTO Statement

We have invested heavily in the current infrastructure, but much of the equipment is approaching the end of its useful life. We are consistently waiting weeks for new gear to be racked before we can start new projects. Our traffic patterns are highest in the mornings and weekend evenings; during other times, 80% of our capacity is sitting idle.

CFO Statement

Our capital expenditure is now exceeding our quarterly projections. Migrating to the cloud will likely cause an initial increase in spending, but we expect to fully transition before our next hardware refresh cycle. Our total cost of ownership (TCO) analysis over the next 5 years puts a cloud strategy between 30 to 50% lower than our current model.

For this question, refer to the Dress4Win case study.

Dress4Win has asked you to recommend machine types they should deploy their application servers to. How should you proceed?
* Perform a mapping of the on-premises physical hardware cores and RAM to the nearest machine types in the cloud.
* Recommend that Dress4Win deploy application servers to machine types that offer the highest RAM to CPU ratio available.
* Recommend that Dress4Win deploy into production with the smallest instances available, monitor them over time, and scale the machine type up until the desired performance is reached.
* Identify the number of virtual cores and RAM associated with the application server virtual machines align them to a custom machine type in the cloud, monitor performance, and scale the machine types up until the desired performance is reached.

Q121. Your company has multiple on-premises systems that serve as sources for reporting. The data has not been maintained well and has become degraded over time. You want to use Google-recommended practices to detect anomalies in your company data. What should you do?
* Upload your files into Cloud Storage. Use Cloud Datalab to explore and clean your data.
* Upload your files into Cloud Storage. Use Cloud Dataprep to explore and clean your data.
* Connect Cloud Datalab to your on-premises systems. Use Cloud Datalab to explore and clean your data.
* Connect Cloud Dataprep to your on-premises systems. Use Cloud Dataprep to explore and clean your data.
https://cloud.google.com/dataprep/

Q122. 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?
* The effective policy is determined only by the policy set at the node
* The effective policy is the policy set at the node and restricted by the policies of its ancestors
* The effective policy is the union of the policy set at the node and policies inherited from its ancestors
* The effective policy is the intersection of the policy set at the node and policies inherited from its ancestors
Reference: https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy

Q123. 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?
* Direct them to download and install the Google StackDriver logging agent.
* Send them a list of online resources about logging best practices.

* Help them define their requirements and assess viable logging tools.
* Help them upgrade their current tool to take advantage of any new features.

Q124. Case Study: 2 – TerramEarth Case Study

Company Overview

TerramEarth manufactures heavy equipment for the mining and agricultural industries: About

80% of their business is from mining and 20% from agriculture. They currently have over 500 dealers and service centers in 100 countries. Their mission is to build products that make their customers more productive.

Company Background

TerramEarth formed in 1946, when several small, family owned companies combined to retool after World War II. The company cares about their employees and customers and considers them to be extended members of their family.

TerramEarth is proud of their ability to innovate on their core products and find new markets as their customers’ needs change. For the past 20 years trends in the industry have been largely toward increasing productivity by using larger vehicles with a human operator.

Solution Concept

There are 20 million TerramEarth vehicles in operation that collect 120 fields of data per second.

Data is stored locally on the vehicle and can be accessed for analysis when a vehicle is serviced.

The data is downloaded via a maintenance port. This same port can be used to adjust operational parameters, allowing the vehicles to be upgraded in the field with new computing modules.

Approximately 200,000 vehicles are connected to a cellular network, allowing TerramEarth to collect data directly. At a rate of 120 fields of data per second, with 22 hours of operation per day.

TerramEarth collects a total of about 9 TB/day from these connected vehicles.

Existing Technical Environment

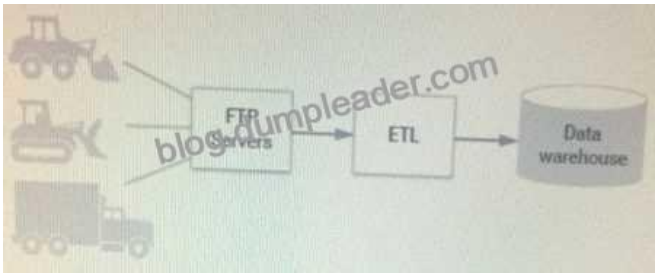

TerramEarth’s existing architecture is composed of Linux-based systems that reside in a data center. These systems gzip CSV files from the field and upload via FTP, transform and aggregate them, and place the data in their data warehouse. Because this

process takes time, aggregated reports are based on data that is 3 weeks old.

With this data, TerramEarth has been able to preemptively stock replacement parts and reduce unplanned downtime of their vehicles by 60%. However, because the data is stale, some customers are without their vehicles for up to 4 weeks while they wait for replacement parts.

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 IP 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

CEO Statement

We have been successful in capitalizing on the trend toward larger vehicles to increase the productivity of our customers. Technological change is occurring rapidly and TerramEarth has taken advantage of connected devices technology to provide our customers with better services, such as our intelligent farming equipment. With this technology, we have been able to increase farmers’ yields by 25%, by using past trends to adjust how our vehicles operate. These advances have led to the rapid growth of our agricultural product line, which we expect will generate 50% of our revenues by 2020.

CTO Statement

Our competitive advantage has always been in the manufacturing process with our ability to build better vehicles for tower cost than our competitors. However, new products with different approaches are constantly being developed, and I’m concerned that we lack the skills to undergo the next wave of transformations in our industry. Unfortunately, our CEO doesn’t take technology obsolescence seriously and he considers the many new companies in our industry to be niche players. My goals are to build our skills while addressing immediate market needs through incremental innovations.

For this question, refer to the TerramEarth case study You analyzed TerramEarth’s business requirement to reduce downtime, and found that they can achieve a majority of time saving by reducing customers’ wait time for parts You decided to focus on reduction of the 3 weeks aggregate reporting time Which modifications to the company’s processes should you recommend?
* Migrate from CSV to binary format, migrate from FTP to SFTP transport, and develop machine learning analysis of metrics.
* Migrate from FTP to streaming transport, migrate from CSV to binary format, and develop machine learning analysis of metrics.
* Increase fleet cellular connectivity to 80%, migrate from FTP to streaming transport, and develop machine learning analysis of metrics.
* Migrate from FTP to SFTP transport, develop machine learning analysis of metrics, and increase dealer local inventory by a fixed factor.
The Avro binary format is the preferred format for loading compressed data. Avro data is faster to load because the data can be read in parallel, even when the data blocks are compressed.

Cloud Storage supports streaming transfers with the gsutil tool or boto library, based on HTTP chunked transfer encoding. Streaming data lets you stream data to and from your Cloud Storage account as soon as it becomes available without requiring that the data be first saved to a separate file. Streaming transfers are useful if you have a process that generates data and you do not want to buffer it locally before uploading it, or if you want to send the result from a computational pipeline directly into Cloud Storage.

References: https://cloud.google.com/storage/docs/streaming

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

Q125. Your company is migrating its on-premises data center into the cloud. As part of the migration, you want to integrate Kubernetes Engine for workload orchestration. Parts of your architecture must also be PCI DSS- compliant. Which of the following is most accurate?
* App Engine is the only compute platform on GCP that is certified for PCI DSS hosting.
* Kubernetes Engine cannot be used under PCI DSS because it is considered shared hosting.
* Kubernetes Engine and GCP provide the tools you need to build a PCI DSS-compliant environment.
* All Google Cloud services are usable because Google Cloud Platform is certified PCI-compliant.

Q126. A few days after JencoMart migrates the user credentials database to Google Cloud Platform and shuts

down the old server, the new database server stops responding to SSH connections. It is still serving

database requests to the application servers correctly.

What three steps should you take to diagnose the problem? Choose 3 answers.
* Delete the virtual machine (VM) and disks and create a new one
* Delete the instance, attach the disk to a new VM, and investigate
* Take a snapshot of the disk and connect to a new machine to investigate
* Check inbound firewall rules for the network the machine is connected to
* Connect the machine to another network with very simple firewall rules and investigate
* Print the Serial Console output for the instance for troubleshooting, activate the interactive console, and

investigate
Explanation/Reference:

Explanation:

D: Handling “Unable to connect on port 22” error message

Possible causes include:

There is no firewall rule allowing SSH access on the port. SSH access on port 22 is enabled on all

▪

Compute Engine instances by default. If you have disabled access, SSH from the Browser will not

work. If you run sshd on a port other than 22, you need to enable the access to that port with a custom

firewall rule.

The firewall rule allowing SSH access is enabled, but is not configured to allow connections from GCP

.

Console services. Source IP addresses for browser-based SSH sessions are dynamically allocated by

GCP Console and can vary from session to session.

F: Handling “Could not connect, retrying…” error

You can verify that the daemon is running by navigating to the serial console output page and looking for

output lines prefixed with the accounts-from-metadata: string. If you are using a standard image but you do

not see these output prefixes in the serial console output, the daemon might be stopped. Reboot the

instance to restart the daemon.

References:

https://cloud.google.com/compute/docs/ssh-in-browser

https://cloud.google.com/compute/docs/ssh-in-browser

Q127.

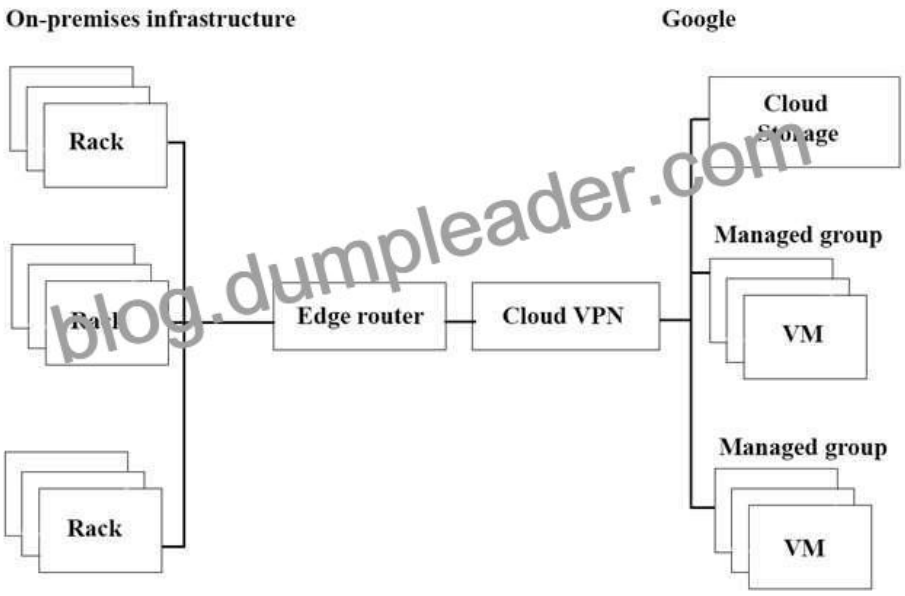

The migration of JencoMart’s application to Google Cloud Platform (GCP) is progressing too slowly. The infrastructure is shown in the diagram. You want to maximize throughput.

What are three potential bottlenecks? (Choose three.)
* A single VPN tunnel, which limits throughput
* A tier of Google Cloud Storage that is not suited for this task
* A copy command that is not suited to operate over long distances
* Fewer virtual machines (VMs) in GCP than on-premises machines
* A separate storage layer outside the VMs, which is not suited for this task
* Complicated internet connectivity between the on-premises infrastructure and GCP

Q128. You need to upgrade the EHR connection to comply with their requirements. The new connection design must support business-critical needs and meet the same network and security policy requirements. What should you do?
* Add a new Dedicated Interconnect connection.
* Upgrade the bandwidth on the Dedicated Interconnect connection to 100 G.
* Add three new Cloud VPN connections.
* Add a new Carrier Peering connection.

Q129. 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
* Port the application code to run on Google App Engine.
* Integrate Cloud Dataflow into the application to capture real-time metrics.
* Instrument the application with a monitoring tool like Stackdriver Debugger.
* Select an automation framework to reliably provision the cloud infrastructure.
* Deploy a continuous integration tool with automated testing in a staging environment.
* Migrate from MySQL to a managed NoSQL database like Google Cloud Datastore or Bigtable.

References: https://cloud.google.com/appengine/docs/standard/java/tools/uploadinganapp

https://cloud.google.com/appengine/docs/standard/java/building-app/cloud-sql