


[Aug-2024 Google Professional-Cloud-Architect Actual Questions and Braindumps [Q114-Q131]



[Aug-2024] Google Professional-Cloud-Architect Actual Questions and Braindumps

Pass Professional-Cloud-Architect Exam with Updated Professional-Cloud-Architect Exam Dumps PDF 2024

Google Professional-Cloud-Architect certification exam is a highly respected certification in the cloud computing industry. It demonstrates the candidate's expertise in designing and managing cloud solutions on GCP and is recognized by employers worldwide. Google Certified Professional - Cloud Architect (GCP) certification also provides access to a network of certified professionals and resources that can help the candidate advance their career in cloud architecture.

Google Professional-Cloud-Architect certification exam is a great way for professionals to validate their skills and knowledge in cloud architecture and design on GCP. Professional-Cloud-Architect exam covers a wide range of topics related to GCP services and requires the candidate to have hands-on experience with GCP tools and services. Google Certified Professional - Cloud Architect (GCP) certification is highly respected in the industry and can help the candidate advance their career in cloud architecture.

QUESTION 114

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 Container Engine, Jenkins, and Helm
- * Google Container Engine and Cloud Load Balancing
- * Google Compute Engine and Cloud Deployment Manager
- * Google Compute Engine, Jenkins, and Cloud Load Balancing

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.

QUESTION 115

Your development teams release new versions of games running on Google Kubernetes Engine (GKE) daily.

You want to create service level indicators (SLIs) to evaluate the quality of the new versions from the user's perspective. What should you do?

- * Create CPU Utilization and Request Latency as service level indicators.
- * Create GKE CPU Utilization and Memory Utilization as service level indicators.
- * Create Server Uptime and Error Rate as service level indicators.
- * Create Request Latency and Error Rate as service level indicators.

QUESTION 116

Your company has an application running on App Engine that allows users to upload music files and share them with other people. You want to allow users to upload files directly into Cloud Storage from their browser session. The payload should not be passed through the backend. What should you do?

1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App Engine application is an allowed origin.
2. Use the Cloud Storage Signed URL feature to generate a POST URL.
 1. Set a CORS configuration in the target Cloud Storage bucket where the base URL of the App Engine application is an allowed

origin.

2. Assign the Cloud Storage WRITER role to users who upload files.

* 1. Use the Cloud Storage Signed URL feature to generate a POST URL.

2. Use App Engine default credentials to sign requests against Cloud Storage.

* 1. Assign the Cloud Storage WRITER role to users who upload files.

2. Use App Engine default credentials to sign requests against Cloud Storage.

QUESTION 117

For this question, refer to the TerraEarth case study. A new architecture that writes all incoming data to BigQuery has been introduced. You notice that the data is dirty, and want to ensure data quality on an automated daily basis while managing cost.

What should you do?

* Set up a streaming Cloud Dataflow job, receiving data by the ingestion process. Clean the data in a Cloud Dataflow pipeline.

* Create a Cloud Function that reads data from BigQuery and cleans it. Trigger it. Trigger the Cloud Function from a Compute Engine instance.

* Create a SQL statement on the data in BigQuery, and save it as a view. Run the view daily, and save the result to a new table.

* Use Cloud Dataprep and configure the BigQuery tables as the source. Schedule a daily job to clean the data.

QUESTION 118

All compute Engine instances in your VPC should be able to connect to an Active Directory server on specific ports. Any other traffic emerging from your instances is not allowed. You want to enforce this using VPC firewall rules.

How should you configure the firewall rules?

* Create an egress rule with priority 1000 to deny all traffic for all instances. Create another egress rule with priority 100 to allow the Active Directory traffic for all instances.

* Create an egress rule with priority 100 to deny all traffic for all instances. Create another egress rule with priority 1000 to allow the Active Directory traffic for all instances.

* Create an egress rule with priority 1000 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 100 to block all traffic for all instances.

* Create an egress rule with priority 100 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 1000 to block all traffic for all instances.

Reference:

<https://cloud.google.com/vpc/docs/firewalls>

QUESTION 119

One of your primary business objectives is being able to trust the data stored in your application. You want to log all changes to the application data. How can you design your logging system to verify authenticity of your logs?

* Write the log concurrently in the cloud and on premises.

* Use a SQL database and limit who can modify the log table.

* Digitally sign each timestamp and log entry and store the signature.

* Create a JSON dump of each log entry and store it in Google Cloud Storage.

Reference:

<https://cloud.google.com/storage/docs/access-logs>

References: <https://cloud.google.com/logging/docs/reference/tools/gcloud-logging>

QUESTION 120

You are migrating your on-premises solution to Google Cloud in several phases. You will use Cloud VPN to maintain a connection between your on-premises systems and Google Cloud until the migration is completed.

You want to make sure all your on-premise systems remain reachable during this period. How should you organize your networking in Google Cloud?

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

QUESTION 121

For this question, refer to the Dress4Win case study. Dress4Win is expected to grow to 10 times its size in 1 year with a corresponding growth in data and traffic that mirrors the existing patterns of usage. The CIO has set the target of migrating production infrastructure to the cloud within the next 6 months. How will you configure the solution to scale for this growth without making major application changes and still maximize the ROI?

- * Migrate the web application layer to App Engine, and MySQL to Cloud Datastore, and NAS to Cloud Storage. Deploy RabbitMQ, and deploy Hadoop servers using Deployment Manager.
- * Migrate RabbitMQ to Cloud Pub/Sub, Hadoop to BigQuery, and NAS to Compute Engine with Persistent Disk storage. Deploy Tomcat, and deploy Nginx using Deployment Manager.
- * Implement managed instance groups for Tomcat and Nginx. Migrate MySQL to Cloud SQL, RabbitMQ to Cloud Pub/Sub, Hadoop to Cloud Dataproc, and NAS to Compute Engine with Persistent Disk storage.
- * Implement managed instance groups for the Tomcat and Nginx. Migrate MySQL to Cloud SQL, RabbitMQ to Cloud Pub/Sub, Hadoop to Cloud Dataproc, and NAS to Cloud Storage.

QUESTION 122

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 collocated 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.

As part of their new application experience, Dress4Win allows customers to upload images of themselves. The customer has exclusive control over who may view these images. Customers should be able to upload images with minimal latency and also be shown their images quickly on the main application page when they log in. Which configuration should Dress4Win use?

- * Store image files in a Google Cloud Storage bucket. Use Google Cloud Datastore to maintain metadata that maps each customer's ID and their image files.
- * Store image files in a Google Cloud Storage bucket. Add custom metadata to the uploaded images in Cloud Storage that contains the customer's unique ID.
- * Use a distributed file system to store customers' images. As storage needs increase, add more persistent disks and/or nodes. Assign each customer a unique ID, which sets each file's owner attribute, ensuring privacy of images.
- * Use a distributed file system to store customers' images. As storage needs increase, add more persistent disks and/or nodes. Use a Google Cloud SQL database to maintain metadata that maps each customer's ID to their image files.

QUESTION 123

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?
 - * Logging, Alerts, Insights, Debug
 - * Monitoring, Trace, Debug, Logging
 - * Monitoring, Logging, Alerts, Error Reporting
 - * Monitoring, Logging, Debug, Error Report

QUESTION 124

Your organization wants to control IAM policies for different departments independently, but centrally.

Which approach should you take?

- * Multiple Organizations with multiple Folders
- * Multiple Organizations, one for each department
- * A single Organization with Folder for each department
- * A single Organization with multiple projects, each with a central owner

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>

QUESTION 125

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?

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

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>

QUESTION 126

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?

- * Grant your colleague the IAM role of project Viewer
- * Perform a rolling restart on the instance group
- * Disable the health check for the instance group. Add his SSH key to the project-wide SSH keys
- * Disable autoscaling for the instance group. Add his SSH key to the project-wide SSH Keys

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

QUESTION 127

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?

- * Google Cloud Storage, Google Cloud Dataflow, Google Compute Engine
- * Google Cloud Storage, Google App Engine, Google Network Load Balancer
- * Google Kubernetes Registry, Google Container Engine, Google HTTP(S) Load Balancer
- * Google Cloud Functions, Google Cloud Pub/Sub, Google Cloud Deployment Manager

Topic 1, Mountkirk Games Case Study 1

Company Overview

Mountkirk Games makes online, session-based, multiplayer games for the most popular mobile platforms.

Company Background

Mountkirk Games builds all of their games with some server-side integration and has historically used cloud providers to lease physical servers. A few of their games were more popular than expected, and they had problems scaling their application servers, MySQL databases, and analytics tools.

Mountkirk's current model is to write game statistics to files and send them through an ETL tool that loads them into a centralized MySQL database for reporting.

Solution Concept

Mountkirk Games is building a new game, which they expect to be very popular. They plan to deploy the game's backend on Google Compute Engine so they can capture streaming metrics, run intensive analytics and take advantage of its autoscaling server environment and integrate with a managed NoSQL database.

Technical Requirements

Requirements for Game Backend Platform

1. Dynamically scale up or down based on game activity.
2. Connect to a managed NoSQL database service.
3. Run customized Linx distro.

Requirements for Game Analytics Platform

1. Dynamically scale up or down based on game activity.
2. Process incoming data on the fly directly from the game servers.
3. Process data that arrives late because of slow mobile networks.
4. Allow SQL queries to access at least 10 TB of historical data.
5. Process files that are regularly uploaded by users; mobile devices.
6. Use only fully managed services

CEO Statement

Our last successful game did not scale well with our previous cloud provider, resulting in lower user adoption and affecting the game's reputation. Our investors want more key performance indicators (KPIs) to evaluate the speed and stability of the game, as well as other metrics that provide deeper insight into usage patterns so we can adapt the games to target users.

CTO Statement

Our current technology stack cannot provide the scale we need, so we want to replace MySQL and move to an environment that provides autoscaling, low latency load balancing, and frees us up from managing physical servers.

CFO Statement

We are not capturing enough user demographic data usage metrics, and other KPIs. As a result, we do not engage the right users. We are not confident that our marketing is targeting the right users, and we are not selling enough premium Blast-Ups inside the games, which dramatically impacts our revenue.

QUESTION 128

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?

- * Create a snapshot of the existing disk. Create an instance template from the snapshot. Create an autoscaled managed instance group from the instance template.
- * Create a snapshot of the existing disk. Create a custom image from the snapshot. Create an autoscaled managed instance group

from the custom image.

- * Create a custom image from the existing disk. Create an instance template from the custom image.

Create an autoscaled managed instance group from the instance template.

- * Create an instance template from the existing disk. Create a custom image from the instance template.

Create an autoscaled managed instance group from the custom image.

Explanation

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

QUESTION 129

An application development team believes their current logging tool will not meet their needs for their new cloud-based product. They want a better 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.

QUESTION 130

You are working in a highly secured environment where public Internet access from the Compute Engine

VMs is not allowed. You do not yet have a VPN connection to access an on-premises file server. You need

to install specific software on a Compute Engine instance. How should you install the software?

- * Upload the required installation files to Cloud Storage. Configure the VM on a subnet with a Private

Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to

the VM using gsutil.

- * Upload the required installation files to Cloud Storage and use firewall rules to block all traffic except

the IP address range for Cloud Storage. Download the files to the VM using gsutil.

- * Upload the required installation files to Cloud Source Repositories. Configure the VM on a subnet with

a Private Google Access subnet. Assign only an internal IP address to the VM. Download the

installation files to the VM using gcloud.

- * Upload the required installation files to Cloud Source Repositories and use firewall rules to block all

traffic except the IP address range for Cloud Source Repositories. Download the files to the VM using

gsutil.

QUESTION 131

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?

- * Add each tier to a different subnet.
- * Set up software based firewalls on individual VMs.
- * Add tags to each tier and set up routes to allow the desired traffic flow.
- * Add tags to each tier and set up firewall rules to allow the desired traffic flow.

<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.

Reference:

<https://aws.amazon.com/it/blogs/aws/building-three-tier-architectures-with-security-groups/>

Topic 2, Mountkirk Games

Company Overview

Mountkirk Games makes online, session-based, multiplayer games for the most popular mobile platforms.

Company Background

Mountkirk Games builds all of their games with some server-side integration and has historically used cloud providers to lease physical servers. A few of their games were more popular than expected, and they had problems scaling their application servers, MySQL databases, and analytics tools.

Mountkirk's current model is to write game statistics to files and send them through an ETL tool that loads them into a centralized MySQL database for reporting.

Solution Concept

Mountkirk Games is building a new game, which they expect to be very popular. They plan to deploy the game's backend on Google Compute Engine so they can capture streaming metrics, run intensive analytics and take advantage of its autoscaling server environment and integrate with a managed NoSQL database.

Technical Requirements

Requirements for Game Backend Platform

1. Dynamically scale up or down based on game activity.
2. Connect to a managed NoSQL database service.
3. Run customized Linux distro.

Requirements for Game Analytics Platform

1. Dynamically scale up or down based on game activity.
2. Process incoming data on the fly directly from the game servers.

3. Process data that arrives late because of slow mobile networks.
4. Allow SQL queries to access at least 10 TB of historical data.
5. Process files that are regularly uploaded by users; mobile devices.
6. Use only fully managed services

CEO Statement

Our last successful game did not scale well with our previous cloud provider, resulting in lower user adoption and affecting the game's reputation. Our investors want more key performance indicators (KPIs) to evaluate the speed and stability of the game, as well as other metrics that provide deeper insight into usage patterns so we can adapt the games to target users.

CTO Statement

Our current technology stack cannot provide the scale we need, so we want to replace MySQL and move to an environment that provides autoscaling, low latency load balancing, and frees us up from managing physical servers.

CFO Statement

We are not capturing enough user demographic data usage metrics, and other KPIs. As a result, we do not engage the right users. We are not confident that our marketing is targeting the right users, and we are not selling enough premium Blast-Ups inside the games, which dramatically impacts our revenue.

Latest Professional-Cloud-Architect Pass Guaranteed Exam Dumps with Accurate & Updated Questions:

<https://www.validexam.com/Professional-Cloud-Architect-latest-dumps.html>