

Exam Professional-Cloud-Architect Tests & New Professional-Cloud-Architect Exam Practice



BTW, DOWNLOAD part of Prep4SureReview Professional-Cloud-Architect dumps from Cloud Storage:
<https://drive.google.com/open?id=1vLjG8940S8eaujOHfIMP7xkmoYRfwYOA>

We often receive news feeds and what well-known entrepreneurs have done to young people. The achievements of these entrepreneurs are the goals we strive for and we must value their opinions. And you may don't know that they were also benefited from our Professional-Cloud-Architect study braindumps. We have engaged in this career for over ten years and helped numerous entrepreneurs achieved their Professional-Cloud-Architect certifications toward their success. Just buy our Professional-Cloud-Architect learning materials and you will become a big man as them.

The GCP certification exam covers a wide range of topics related to cloud architecture, including designing and planning a cloud solution architecture, managing and provisioning infrastructure, ensuring security and compliance, and optimizing technical and business processes. Professional-Cloud-Architect exam is designed to assess the candidate's ability to design, develop, and manage scalable, efficient, and secure cloud solutions using GCP technologies.

Google Professional-Cloud-Architect exam is a certification program that focuses on testing an individual's ability to design and manage secure, scalable, and reliable cloud-based solutions using Google Cloud Platform (GCP). Google Certified Professional - Cloud Architect (GCP) certification is aimed at cloud architects, engineers, and consultants who want to showcase their expertise in Google Cloud technologies and help their organizations migrate to the cloud.

The Google Professional-Cloud-Architect Exam is made up of multiple-choice questions that cover a range of topics, including cloud architecture, security, networking, and data management. Candidates are expected to have a strong understanding of cloud computing technologies and best practices, as well as the ability to apply this knowledge to real-world scenarios. Professional-Cloud-Architect exam is designed to be challenging, with a passing score of 70% required to earn the certification.

New Professional-Cloud-Architect Exam Practice & Training Professional-Cloud-Architect Online

You will need to pass the Google Certified Professional - Cloud Architect (GCP) (Professional-Cloud-Architect) exam to achieve the Google Professional-Cloud-Architect certification. Due to extremely high competition, passing the Google Professional-Cloud-Architect exam is not easy; however, possible. You can use Prep4SureReview products to pass the Professional-Cloud-Architect Exam on the first attempt. The Google practice exam gives you confidence and helps you understand the criteria of the testing authority and pass the Google Certified Professional - Cloud Architect (GCP) (Professional-Cloud-Architect) exam on the first attempt.

Google Certified Professional - Cloud Architect (GCP) Sample Questions (Q320-Q325):

NEW QUESTION # 320

The development team has provided you with a Kubernetes Deployment file. You have no infrastructure yet and need to deploy the application. What should you do?

- A. Use kubectl to create a Kubernetes cluster. Use Deployment Manager to create the deployment.
- B. Use kubectl to create a Kubernetes cluster. Use kubectl to create the deployment.
- C. Use gcloud to create a Kubernetes cluster. Use kubectl to create the deployment.
- D. Use gcloud to create a Kubernetes cluster. Use Deployment Manager to create the deployment.

Answer: C

NEW QUESTION # 321

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

Answer: A,E

Explanation:

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 # 322

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 a single G Suite account to manage users with each stage of each application in its own project.
- B. 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.

- C. Create two G Suite accounts to manage users: one for development/test/staging and one for production. Each account should contain one project for every application.
- D. 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>

Topic 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.

NEW QUESTION # 323

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

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

Answer: A

Explanation:

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

NEW QUESTION # 324

Your company has an application running as a Deployment in a Google Kubernetes Engine (GKE) cluster. When releasing new versions of the application via a rolling deployment, the team has been causing outages. The root cause of the outages is misconfigurations with parameters that are only used in production. You want to put preventive measures for this in the platform to prevent outages. What should you do?

- A. Configure an uptime alert in Cloud Monitoring.
- B. Create a Scheduled Task to check whether the application is available.
- **C. Configure health checks on the managed instance group.**
- D. Configure liveness and readiness probes in the Pod specification.

Answer: C

Explanation:

This option can help prevent outages caused by misconfigurations with parameters that are only used in production. Liveness and readiness probes are mechanisms to check the health and availability of the Pods and containers in a GKE cluster. Liveness probes determine if a container is still running, and if not, restart it.

Readiness probes determine if a container is ready to serve requests, and if not, remove it from the load balancer. By configuring liveness and readiness probes in the Pod specification, you can ensure that your application can handle traffic and recover from failures gracefully during a rolling update. The other options are not optimal for this scenario, because they either do not prevent outages, but only alert or monitor them (B, C), or do not apply to GKE clusters, but to Compute Engine instances (D). References:

<https://cloud.google.com/kubernetes-engine/docs/how-to/updates-apps>

<https://cloud.google.com/blog/products/containers-kubernetes/kubernetes-best-practices-setting-up-health>

