

100% Pass Quiz Professional-Cloud-Architect - Valid Google Certified Professional - Cloud Architect (GCP) Reliable Braindumps Questions



BONUS!!! Download part of Actualtests4sure Professional-Cloud-Architect dumps for free: <https://drive.google.com/open?id=1u7W2IqkI8laNNbQ8Mc9TAc34c9bqFrB9>

As one of the most professional dealer of Professional-Cloud-Architect practice questions, we have connection with all academic institutions in this line with proficient researchers of the knowledge related with the Professional-Cloud-Architect exam materials to meet your tastes and needs, please feel free to choose. And we have three versions of Professional-Cloud-Architect training guide: the PDF, Software and APP online for you. You can choose the one which you like best.

To become a Google Certified Professional - Cloud Architect (GCP), a candidate must pass the Professional-Cloud-Architect certification exam. Professional-Cloud-Architect exam consists of multiple-choice and scenario-based questions that test the candidate's ability to design, manage, and secure cloud solutions on GCP. Professional-Cloud-Architect exam is proctored and can be taken online or at a testing center.

To become a certified Professional-Cloud-Architect, candidates must pass an exam that consists of 50 multiple-choice and multiple-select questions, which must be completed within two hours. Professional-Cloud-Architect Exam is designed to test candidates' ability to design and plan a cloud solution architecture, manage and provision the GCP infrastructure, and optimize technical and business processes.

>> **Professional-Cloud-Architect Reliable Braindumps Questions** <<

New Professional-Cloud-Architect Exam Dumps, Professional-Cloud-Architect Exam Cram

Do you want to prove your ability in IT field? Do you want to get more recognition and employment opportunities? So Professional-Cloud-Architect exam certification will be an important evidence to prove yourself. Almost all those who are working in the IT field know how important to get Professional-Cloud-Architect exam certification. As we know, everyone's energy is limited, if you want to pass the important Professional-Cloud-Architect Certification Exam in such short time, the exam software provided by our Actualtests4sure will be a good helper for your preparation for the exam. The complete questions and exam software created in accordance with the laws of the people's memory will help you succeed in the Professional-Cloud-Architect exam.

Exam Overview

The qualifying test for the Google Professional Cloud Architect certification is 2 hours long. The questions administered during the exam can be presented as multiple-select and multiple-choice items. The test is delivered in English or Japanese and can be taken by the students either online from a remote location or in-person at the authorized testing center. To check the location of the nearest center, you can use the search on the official website. To schedule the exam, each applicant is required to pay the registration fee of \$200, plus applicable taxes.

Google Certified Professional - Cloud Architect (GCP) Sample Questions (Q343-Q348):

NEW QUESTION # 343

For this question, refer to the TerramEarth case study. You need to implement a reliable, scalable GCP solution for the data warehouse for your company, TerramEarth. Considering the TerramEarth business and technical requirements, what should you do?

- A. Replace the existing data warehouse with BigQuery. Use table partitioning.
- **B. Replace the existing data warehouse with BigQuery. Use federated data sources.**
- C. Replace the existing data warehouse with a Compute Engine instance with 96 CPUs. Add an additional Compute Engine pre-emptible instance with 32 CPUs.
- D. Replace the existing data warehouse with a Compute Engine instance with 96 CPUs.

Answer: B

Explanation:

Reference:

https://cloud.google.com/solutions/bigquery-data-warehouse#external_sources

<https://cloud.google.com/solutions/bigquery-data-warehouse>

NEW QUESTION # 344

Case Study: 6 - TerramEarth

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.

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 and Windows-based systems that reside in a single U.S. west coast based data center. These systems gzip CSV files from the field and upload via FTP, 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.
- Support the dealer network with more data on how their customers use their equipment to better position new products and services
- Have the ability to partner with different companies - especially with seed and fertilizer suppliers in the fast-growing agricultural business - to create compelling joint offerings for their customers.

Technical Requirements

- Expand beyond a single datacenter to decrease latency to the American Midwest and east coast.
- Create a backup strategy.
- Increase security of data transfer from equipment to the datacenter.
- Improve data in the data warehouse.
- Use customer and equipment data to anticipate customer needs.

Application 1: Data ingest

A custom Python application reads uploaded datafiles from a single server, writes to the data warehouse.

Compute:

- Windows Server 2008 R2
- 16 CPUs

- 128 GB of RAM
- 10 TB local HDD storage

Application 2: Reporting

An off the shelf application that business analysts use to run a daily report to see what equipment needs repair. Only 2 analysts of a team of 10 (5 west coast, 5 east coast) can connect to the reporting application at a time.

Compute:

- Off the shelf application. License tied to number of physical CPUs
- Windows Server 2008 R2
- 16 CPUs
- 32 GB of RAM
- 500 GB HDD

Data warehouse:

- A single PostgreSQL server
- RedHat Linux
- 64 CPUs
- 128 GB of RAM
- 4x 6TB HDD in RAID 0

Executive Statement

Our competitive advantage has always been in the manufacturing process, with our ability to build better vehicles for lower 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. My goals are to build our skills while addressing immediate market needs through incremental innovations.

You need to implement a reliable, scalable GCP solution for the data warehouse for your company, TerramEarth. Considering the TerramEarth business and technical requirements, what should you do?

- A. Replace the existing data warehouse with BigQuery. Use federated data sources.
- **B. Replace the existing data warehouse with BigQuery. Use table partitioning.**
- C. Replace the existing data warehouse with a Compute Engine instance with 96 CPUs. Add an additional Compute Engine pre-emptible instance with 32 CPUs.
- D. Replace the existing data warehouse with a Compute Engine instance with 96 CPUs.

Answer: B

Explanation:

1. BigQuery does not guarantee data consistency for external data sources. Changes to the underlying data while a query is running can result in unexpected behavior.
2. Query performance for external data sources may not be as high as querying data in a native BigQuery table.

NEW QUESTION # 345

Your customer is moving an existing corporate application to Google Cloud Platform from an on- premises data center. The business owners require minimal user disruption. There are strict security team requirements for storing passwords. What authentication strategy should they use?

- **A. Provision users in Google using the Google Cloud Directory Sync tool.**
- B. Use G Suite Password Sync to replicate passwords into Google.
- C. Federate authentication via SAML 2.0 to the existing Identity Provider.
- D. Ask users to set their Google password to match their corporate password.

Answer: A

Explanation:

Provision users to Google's directory

The global Directory is available to both Cloud Platform and G Suite resources and can be provisioned by a number of means. Provisioned users can take advantage of rich authentication features including single sign-on (SSO), OAuth, and two-factor verification.

You can provision users automatically using one of the following tools and services:

Google Cloud Directory Sync (GCDS)

Google Admin SDK

A third-party connector

GCDS is a connector that can provision users and groups on your behalf for both Cloud Platform and G Suite. Using GCDS, you

can automate the addition, modification, and deletion of users, groups, and non-employee contacts. You can synchronize the data from your LDAP directory server to your Cloud Platform domain by using LDAP queries. This synchronization is one-way: the data in your LDAP directory server is never modified.

References: <https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#authentication-and-identity>

NEW QUESTION # 346

You are deploying a new three-tier application to Compute Engine instances within a single Virtual Private Cloud (VPC). The architecture is segmented into three subnets: a web tier subnet, an application tier subnet, and a database tier subnet.

- The web tier must only receive traffic from an external load balancer.
- The application tier must only receive traffic from the web tier.
- The database tier must only receive traffic from the application tier.

You need to enforce strict traffic flow control and want to follow Google-recommended practices.

What should you do?

- A. Set up VPC Network Peering between the web tier and application tier subnets, and another peering connection between the application tier and database tier subnets.
- B. Configure Cloud NAT for each subnet, and create Google Cloud Armor policies to filter traffic between the tiers based on their IP address ranges.
- **C. Use a combination of network tags and service accounts. Apply a unique network tag and a dedicated service account to the instances in each tier. Then create specific firewall rules that allow ingress traffic based on the source service account or tag of the upstream tier.**
- D. Create a single VPC firewall rule with a high priority that allows ingress traffic on all ports between the IP address ranges of the three subnets.

Answer: C

Explanation:

Using network tags and/or service accounts to scope firewall rules is the recommended way to enforce least-privilege traffic flows within a VPC. By assigning distinct tags and service accounts per tier, you can create precise ingress rules that only allow traffic from the upstream tier and block all other sources, ensuring the web tier only accepts load balancer traffic, the app tier only accepts web-tier traffic, and the database tier only accepts app-tier traffic.

NEW QUESTION # 347

You have an application that will run on Compute Engine. You need to design an architecture that takes into account a disaster recovery plan that requires your application to fail over to another region in case of a regional outage. What should you do?

- A. Deploy the application on two Compute Engine instance groups, each in the same project but in a different region. Use the first instance group to serve traffic, and use the HTTP load balancing service to fail over to the standby instance group in case of a disaster.
- B. Deploy the application on two Compute Engine instance groups, each in separate project and a different region. Use the first instance group to server traffic, and use the HTTP load balancing service to fail over to the standby instance in case of a disaster.
- **C. Deploy the application on a Compute Engine instance. Use the instance to serve traffic, and use the HTTP load balancing service to fail over to an instance on your premises in case of a disaster.**
- D. Deploy the application on two Compute Engine instances in the same project but in a different region. Use the first instance to serve traffic, and use the HTTP load balancing service to fail over to the standby instance in case of a disaster.

Answer: C

NEW QUESTION # 348

.....

New Professional-Cloud-Architect Exam Dumps: <https://www.actualtests4sure.com/Professional-Cloud-Architect-test-questions.html>

