# Q&As

Professional Cloud Architect on Google Cloud Platform

# Pass Google PROFESSIONAL-CLOUD-ARCHITECT Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.geekcert.com/professional-cloud-architect.html

## 100% Passing Guarantee
## 100% Money Back Assurance

Following Questions and Answers are all new published by Google Official Exam Center

PROFESSIONAL-CLOUD-ARCHITECT PDF Dumps | PROFESSIONAL-CLOUD-ARCHITECT Study Guide | PROFESSIONAL-CLOUD-ARCHITECT Braindumps

1 / 5

**QUESTION 1**

You are working at an institution that processes medical data. You are migrating several workloads onto Google Cloud. Company policies require all workloads to run on physically separated hardware, and workloads from different clients must also be separated You created a sole-tenant node group and added a node for each client. You need to deploy the workloads on these dedicated hosts. What should you do?

A. Add the node group name as a network tag when creating Compute Engine instances in order to host each workload on the correct node group.

B. Add the node name as a network tag when creating Compute Engine instances in order to host each workload on the correct node.

C. Use node affinity labels based on the node group name when creating Compute Engine instances in order to host each workload on the correct node group

D. Use node affinity labels based on the node name when creating Compute Engine instances in order to host each workload on the correct node.

Correct Answer: C

https://cloud.google.com/compute/docs/nodes/provisioning-sole-tenant-vms#provision_a_sole-tenant_vm
https://cloud.google.com/compute/docs/nodes/provisioning-sole-tenant-vms#gcloud_2 When you create a VM, you request sole-tenancy by specifying node affinity or anti-affinity, referencing one or more node affinity labels. You specify custom node affinity labels when you create a node template, and Compute Engine automatically includes some default affinity labels on each node. By specifying affinity when you create a VM, you can schedule VMs together on a specific node or nodes in a node group. By specifying anti-affinity when you create a VM, you can ensure that certain VMs are not scheduled together on the same node or nodes in a node group.

---

**QUESTION 2**

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 transfer archived data to Cloud Storage. Establish a connection with Google using a Dedicated Interconnect or Direct Peering connection and use it to upload files daily.

C. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transfer archived data to Cloud Storage. Establish one Cloud VPN Tunnel to VPC networks over the public internet, and compares and upload files daily using the gsutil -m option.

D. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storage. Establish a Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily.

Correct Answer: B

---

https://cloud.google.com/interconnect/docs/how-to/direct-peering

**QUESTION 3**

You are designing a Data Warehouse on Google Cloud and want to store sensitive data in BigQuery. Your company requires you to generate encryption keys outside of Google Cloud. You need to implement a solution. What should you do?

A. Generate a new key in Cloud Key Management Service (Cloud KMS). Store all data in Cloud Storage using the customer-managed key option and select the created key. Set up a Dataflow pipeline to decrypt the data and to store it in a BigQuery dataset.

B. Generate a new key in Cloud Key Management Service (Cloud KMS). Create a dataset in BigQuery using the customer-managed key option and select the created key

C. Import a key in Cloud KMS. Store all data in Cloud Storage using the customer-managed key option and select the created key. Set up a Dataflow pipeline to decrypt the data and to store it in a new BigQuery dataset.

D. Import a key in Cloud KMS. Create a dataset in BigQuery using the customer-supplied key option and select the created key.

Correct Answer: D

https://cloud.google.com/bigquery/docs/customer-managed-encryption

**QUESTION 4**

A lead engineer wrote a custom tool that deploys virtual machines in the legacy data center.

He wants to migrate the custom tool to the new cloud environment

You want to advocate for the adoption of Google Cloud Deployment Manager

What are two business risks of migrating to Cloud Deployment Manager? Choose 2 answers

A. Cloud Deployment Manager uses Python.

B. Cloud Deployment Manager APIs could be deprecated in the future.

C. Cloud Deployment Manager is unfamiliar to the company\\'s engineers.

D. Cloud Deployment Manager requires a Google APIs service account to run.

E. Cloud Deployment Manager can be used to permanently delete cloud resources.

F. Cloud Deployment Manager only supports automation of Google Cloud resources.

Correct Answer: CF

https://cloud.google.com/deployment-manager/docs/deployments/deleting-deployments

**QUESTION 5**

Your applications will be writing their logs to BigQuery for analysis. Each application should have its own table.

Any logs older than 45 days should be removed. You want to optimize storage and follow Google recommended practices. What should you do?

A. Configure the expiration time for your tables at 45 days

B. Make the tables time-partitioned, and configure the partition expiration at 45 days

C. Rely on BigQuery\\'s default behavior to prune application logs older than 45 days

D. Create a script that uses the BigQuery command line tool (bq) to remove records older than 45 days

Correct Answer: B

https://cloud.google.com/bigquery/docs/managing-partitioned-tables

[PROFESSIONAL-CLOUD-ARCHITECT PDF Dumps](#)

[PROFESSIONAL-CLOUD-ARCHITECT Study Guide](#)

[PROFESSIONAL-CLOUD-ARCHITECT Braindumps](#)

[PROFESSIONAL-CLOUD-ARCHITECT PDF Dumps](#) | [PROFESSIONAL-CLOUD-ARCHITECT Study Guide](#) | [PROFESSIONAL-CLOUD-ARCHITECT Braindumps](#)

5 / 5