

- A. Add a new Dedicated Interconnect connection
- B. Upgrade the bandwidth on the Dedicated Interconnect connection to 100 G
- C. Add three new Cloud VPN connections
- D. Add a new Carrier Peering connection

**Correct Answer: A**

**Explanation:**

The case does not call out the throughput being an issue. However, to achieve 99.99%, you need to have 4 connections as per Google recommendations. However, in the options only A has the option to add an additional Interconnect connection.

<https://cloud.google.com/network-connectivity/docs/interconnect/concepts/dedicated-overview#availability>

## Topic 10, Misc Questions

### QUESTION 1

Your company is planning to perform a lift and shift migration of their Linux RHEL 6.5+ virtual machines. The virtual machines are running in an on-premises VMware environment. You want to migrate them to Compute Engine following Google-recommended practices. What should you do?

- A.
  - 1. Define a migration plan based on the list of the applications and their dependencies.
  - 2. Migrate all virtual machines into Compute Engine individually with Migrate for Compute Engine.
- B.
  - 1. Perform an assessment of virtual machines running in the current VMware environment.
  - 2. Create images of all disks. Import disks on Compute Engine.
  - 3. Create standard virtual machines where the boot disks are the ones you have imported.
- C.
  - 1. Perform an assessment of virtual machines running in the current VMware environment.
  - 2. Define a migration plan, prepare a Migrate for Compute Engine migration RunBook, and execute the migration.
- D.
  - 1. Perform an assessment of virtual machines running in the current VMware environment.
  - 2. Install a third-party agent on all selected virtual machines.
  - 3. Migrate all virtual machines into Compute Engine.

**Correct Answer: C**

**Explanation:**

The framework illustrated in the preceding diagram has four phases:

**Assess.** In this phase, you assess your source environment, assess the workloads that you want to migrate to Google Cloud, and assess which VMs support each workload.

**Plan.** In this phase, you create the basic infrastructure for Migrate for Compute Engine, such as provisioning the resource hierarchy and setting up network access.

**Deploy.** In this phase, you migrate the VMs from the source environment to Compute Engine.

**Optimize.** In this phase, you begin to take advantage of the cloud technologies and capabilities.

Reference: <https://cloud.google.com/architecture/migrating-vms-migrate-for-compute-engine-getting-started>

### QUESTION 2

Your company is using BigQuery as its enterprise data warehouse. Data is distributed over several Google Cloud projects. All queries on BigQuery need to be billed on a single project. You want to make sure that no query costs are incurred on the projects that contain the data. Users should be able to query the datasets, but not edit them. How should you configure users' access roles?

[Professional-Cloud-Architect Exam Dumps](#)   [Professional-Cloud-Architect PDF Dumps](#)

[Professional-Cloud-Architect VCE Dumps](#)   [Professional-Cloud-Architect Q&As](#)

<https://www.ensurepass.com/PCA.html>

- A. Add all users to a group. Grant the group the role of BigQuery user on the billing project and BigQuery dataViewer on the projects that contain the data.
- B. Add all users to a group. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery user on the projects that contain the data.
- C. Add all users to a group. Grant the group the roles of BigQuery jobUser on the billing project and BigQuery dataViewer on the projects that contain the data.
- D. Add all users to a group. Grant the group the roles of BigQuery dataViewer on the billing project and BigQuery jobUser on the projects that contain the data.

**Correct Answer:** A

**Explanation:**

<https://cloud.google.com/bigquery/docs/running-queries>

### QUESTION 3

Your company recently acquired a company that has infrastructure in Google Cloud. Each company has its own Google Cloud organization. Each company is using a Shared Virtual Private Cloud (VPC) to provide network connectivity for its applications. Some of the subnets used by both companies overlap. In order for both businesses to integrate, the applications need to have private network connectivity. These applications are not on overlapping subnets. You want to provide connectivity with minimal re-engineering. What should you do?

- A. Set up VPC peering and peer each Shared VPC together
- B. Configure SSH port forwarding on each application to provide connectivity between applications in the different Shared VPCs
- C. Migrate the protects from the acquired company into your company's Google Cloud organization. Re launch the instances in your companies Shared VPC
- D. Set up a Cloud VPN gateway in each Shared VPC and peer Cloud VPNs

**Correct Answer:** B

### QUESTION 4

Your company has developed a monolithic, 3-tier application to allow external users to upload and share files. The solution cannot be easily enhanced and lacks reliability. The development team would like to re-architect the application to adopt microservices and a fully managed service approach, but they need to convince their leadership that the effort is worthwhile. Which advantage(s) should they highlight to leadership?

- A. The new approach will be significantly less costly, make it easier to manage the underlying infrastructure, and automatically manage the CI/CD pipelines.
- B. The monolithic solution can be converted to a container with Docker. The generated container can then be deployed into a Kubernetes cluster.
- C. The new approach will make it easier to decouple infrastructure from application, develop and release new features, manage the underlying infrastructure, manage CI/CD pipelines and perform A/B testing, and scale the solution if necessary.
- D. The process can be automated with Migrate for Compute Engine.

**Correct Answer:** C

**Explanation:**

The new approach will make it easier to decouple infrastructure from an application, develop and release new features, manage the underlying infrastructure, manage CI/CD pipelines and perform A/B testing, and scale the solution if necessary.

[Professional-Cloud-Architect Exam Dumps](#)   [Professional-Cloud-Architect PDF Dumps](#)

[Professional-Cloud-Architect VCE Dumps](#)   [Professional-Cloud-Architect Q&As](#)

<https://www.ensurepass.com/PCA.html>

#### QUESTION 5

Your company has an application running on Compute Engine that allows users to play their favorite music. There are a fixed number of instances. Files are stored in Cloud Storage and data is streamed directly to users. Users are reporting that they sometimes need to attempt to play popular songs multiple times before they are successful. You need to improve the performance of the application. What should you do?

- A. 1. Copy popular songs into CloudSQL as a blob  
2. Update application code to retrieve data from CloudSQL when Cloud Storage is overloaded
- B. 1. Create a managed instance group with Compute Engine instances  
2. Create a global load balancer and configure it with two backends
  - Managed instance group
  - Cloud Storage bucket
- C. 1. Enable Cloud CDN on the bucket backend
- D. 1. Mount the Cloud Storage bucket using gcsfuse on all backend Compute Engine instances  
2. Serve music files directly from the backend Compute Engine instance
- E. 1. Create a Cloud Filestore NFS volume and attach it to the backend Compute Engine instances  
2. Download popular songs in Cloud Filestore  
3. Serve music files directly from the backend Compute Engine instance

**Correct Answer: B**

#### QUESTION 6

Your company has just recently activated Cloud Identity to manage users. The Google Cloud Organization has been configured as usual. The security team needs to secure permissions that will be part of the Organization. They want to prohibit IAM users outside the domain from gaining permissions from now on. What should they do?

- A. Configure an organization policy to restrict identities by domain
- B. Configure an organization policy to block creation of service accounts
- C. Configure Cloud Scheduler to trigger a Cloud Function every hour that removes all users that don't belong to the Cloud Identity domain from all projects.
- D. None of the above

**Correct Answer: A**

#### QUESTION 7

Your company is moving 75 TB of data into Google Cloud. You want to use Cloud Storage and follow Google-recommended practices. What should you do?

- A. Move your data onto a Transfer Appliance. Use a Transfer Appliance Rehydrator to decrypt the data into Cloud Storage.
- B. Move your data onto a Transfer Appliance. Use Cloud Dataprep to decrypt the data into Cloud Storage.
- C. Install gsutil on each server that contains data. Use resumable transfers to upload the data into Cloud Storage.
- D. Install gsutil on each server containing data. Use streaming transfers to upload the data into Cloud Storage.

**Correct Answer: A**

**Explanation:**

<https://cloud.google.com/transfer-appliance/docs/2.0/faq>

**QUESTION 8**

You are developing an application using different microservices that should remain internal to the cluster. You want to be able to configure each microservice with a specific number of replicas. You also want to be able to address a specific microservice from any other microservice in a uniform way, regardless of the number of replicas the microservice scales to. You need to implement this solution on Google Kubernetes Engine. What should you do?

- A. Deploy each microservice as a Deployment. Expose the Deployment in the cluster using a Service, and use the Service DNS name to address it from other microservices within the cluster.
- B. Deploy each microservice as a Deployment. Expose the Deployment in the cluster using an Ingress, and use the Ingress IP address to address the Deployment from other microservices within the cluster.
- C. Deploy each microservice as a Pod. Expose the Pod in the cluster using a Service, and use the Service DNS name to address the microservice from other microservices within the cluster.
- D. Deploy each microservice as a Pod. Expose the Pod in the cluster using an Ingress, and use the Ingress IP address name to address the Pod from other microservices within the cluster.

**Correct Answer: A**

**Explanation:**

<https://kubernetes.io/docs/concepts/services-networking/ingress/>

**QUESTION 9**

A news feed web service has the following code running on Google App Engine. During peak load, users report that they can see news articles they already viewed. What is the most likely cause of this problem?

```
import news
from flask import Flask, redirect, request
from flask.ext.api import status
from google.appengine.api import users

app = Flask(__name__)
sessions = {}

@app.route("/")
def homepage():
    user = users.get_current_user()
    if not user:
        return "Invalid login",
        status.HTTP_401_UNAUTHORIZED

    if user not in sessions:
        sessions[user] = {"viewed": []}

    news_articles = news.get_new_news (user, sessions [user]
["viewed"])
    sessions [user] ["viewed"] += [n["id"] for n
in news_articles]

    return news.render(news_articles)

if __name__ == "__main__":
    app.run()
```

- A. The session variable is local to just a single instance.

[Professional-Cloud-Architect Exam Dumps](#)   [Professional-Cloud-Architect PDF Dumps](#)

[Professional-Cloud-Architect VCE Dumps](#)   [Professional-Cloud-Architect Q&As](#)

<https://www.ensurepass.com/PCA.html>

- B. The session variable is being overwritten in Cloud Datastore.
- C. The URL of the API needs to be modified to prevent caching.
- D. The HTTP Expires header needs to be set to -1 to stop caching.

**Correct Answer: A**

**Explanation:**

<https://stackoverflow.com/questions/3164280/google-app-engine-cache-list-in-session-variable?rq=1>

#### QUESTION 10

You company has a Kubernetes application that pulls messages from Pub/Sub and stores them in Firestore. Because the application is simple, it was deployed as a single pod. The infrastructure team has analyzed Pub/Sub metrics and discovered that the application cannot process the messages in real time. Most of them wait for minutes before being processed. You need to scale the elaboration process that is I/O-intensive. What should you do?

- A. Configure a Kubernetes autoscaling based on the subscription/push\_request metric.
- B. Use the -enable -autoscaling flag when you create the Kubernetes cluster
- C. Configure a Kubernetes autoscaling based on the subscription/num\_undelivered message metric.
- D. Use kubectl autoscale deployment APP\_NAME -max 6 -min 2 -cpu- percent 50 to configure Kubernetes autoscaling deployment

**Correct Answer: A**

**Explanation:**

[https://cloud.google.com/kubernetes-engine/docs/concepts/custom-and-external-metrics#external\\_metrics](https://cloud.google.com/kubernetes-engine/docs/concepts/custom-and-external-metrics#external_metrics)

#### QUESTION 11

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

**Correct Answer: B**

**Explanation:**

<https://cloud.google.com/kubernetes-engine/docs/how-to/creating-a-cluster>

#### QUESTION 12

You are designing an application for use only during business hours. For the minimum viable product release, you'd like to use a managed product that automatically "scales to zero" so you don't incur costs when there is no activity. Which primary compute resource should you choose?

- A. Cloud Functions
- B. Compute Engine
- C. Kubernetes Engine
- D. AppEngine flexible environment