



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.passapply.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



VCE & PDF

PassApply.com

<https://www.passapply.com/professional-cloud-architect.html>

2024 Latest passapply PROFESSIONAL-CLOUD-ARCHITECT PDF and VCE dumps Download

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers





QUESTION 1

You need to design a solution for global load balancing based on the URL path being requested. You need to ensure operations reliability and end-to-end in-transit encryption based on Google best practices. What should you do?

- A. Create a cross-region load balancer with URL Maps.
- B. Create an HTTPS load balancer with URL maps.
- C. Create appropriate instance groups and instances. Configure SSL proxy load balancing.
- D. Create a global forwarding rule. Configure SSL proxy balancing.

Correct Answer: B

Reference <https://cloud.google.com/load-balancing/docs/https/url-map>

QUESTION 2

Your organization wants to control IAM policies for different departments independently, but centrally. Which approach should you take?

- A. Multiple Organizations with multiple Folders
- B. Multiple Organizations, one for each department
- C. A single Organization with Folders for each department
- D. A single Organization with multiple projects, each with a central owner

Correct Answer: C

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 3

For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the compute workloads for your company, Mountkirk Games. Considering the Mountkirk games business and technical requirements, what should you do?

- A. Create network load balancers. Use preemptible Compute Engine instances.
- B. Create network load balancers. Use non-preemptible Compute Engine instances.
- C. Create a global load balancer with managed instance groups and autoscaling policies. Use preemptible Compute



Engine instances.

D. Create a global load balancer with managed instance groups and autoscaling policies. Use non-preemptible Compute Engine instance.

Correct Answer: D

QUESTION 4

Your organization has decided to restrict the use of external IP addresses on instances to only approved instances. You want to enforce this requirement across all of your Virtual Private Clouds (VPCs). What should you do?

A. Remove the default route on all VPCs. Move all approved instances into a new subnet that has a default route to an internet gateway.

B. Create a new VPC in custom mode. Create a new subnet for the approved instances, and set a default route to the internet gateway on this new subnet.

C. Implement a Cloud NAT solution to remove the need for external IP addresses entirely.

D. Set an Organization Policy with a constraint on constraints/compute.vmExternalIpAccess. List the approved instances in the allowedValues list.

Correct Answer: D

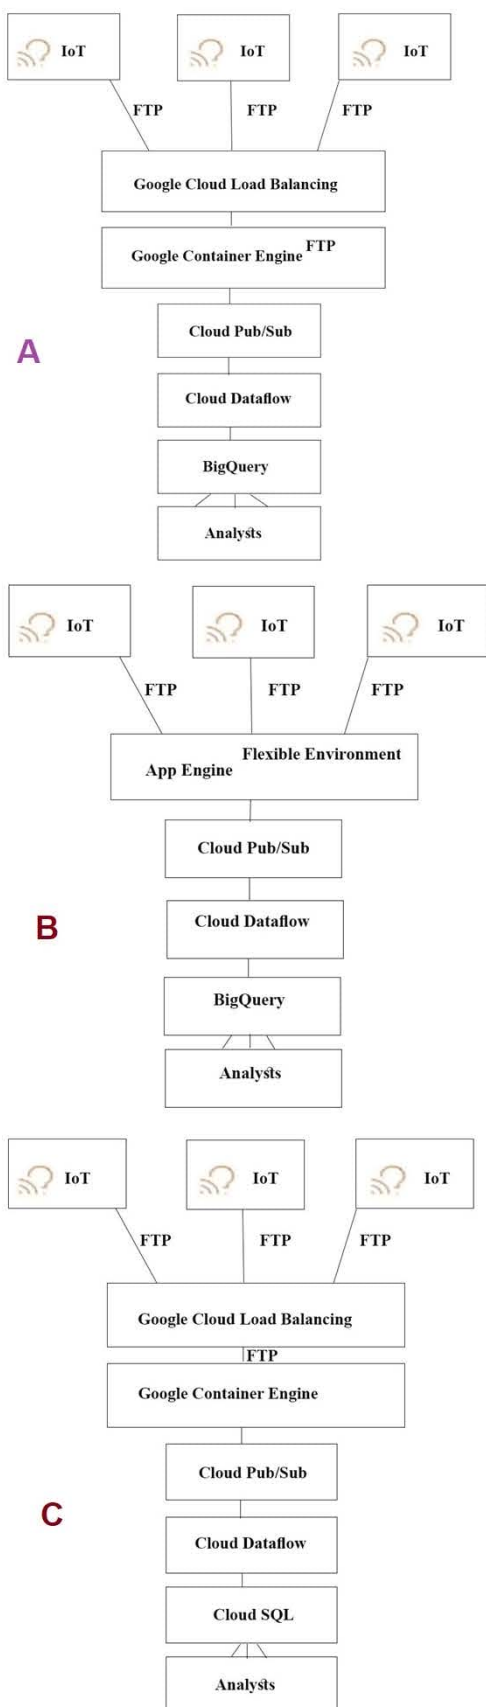
Reference: <https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#disableexternalip>

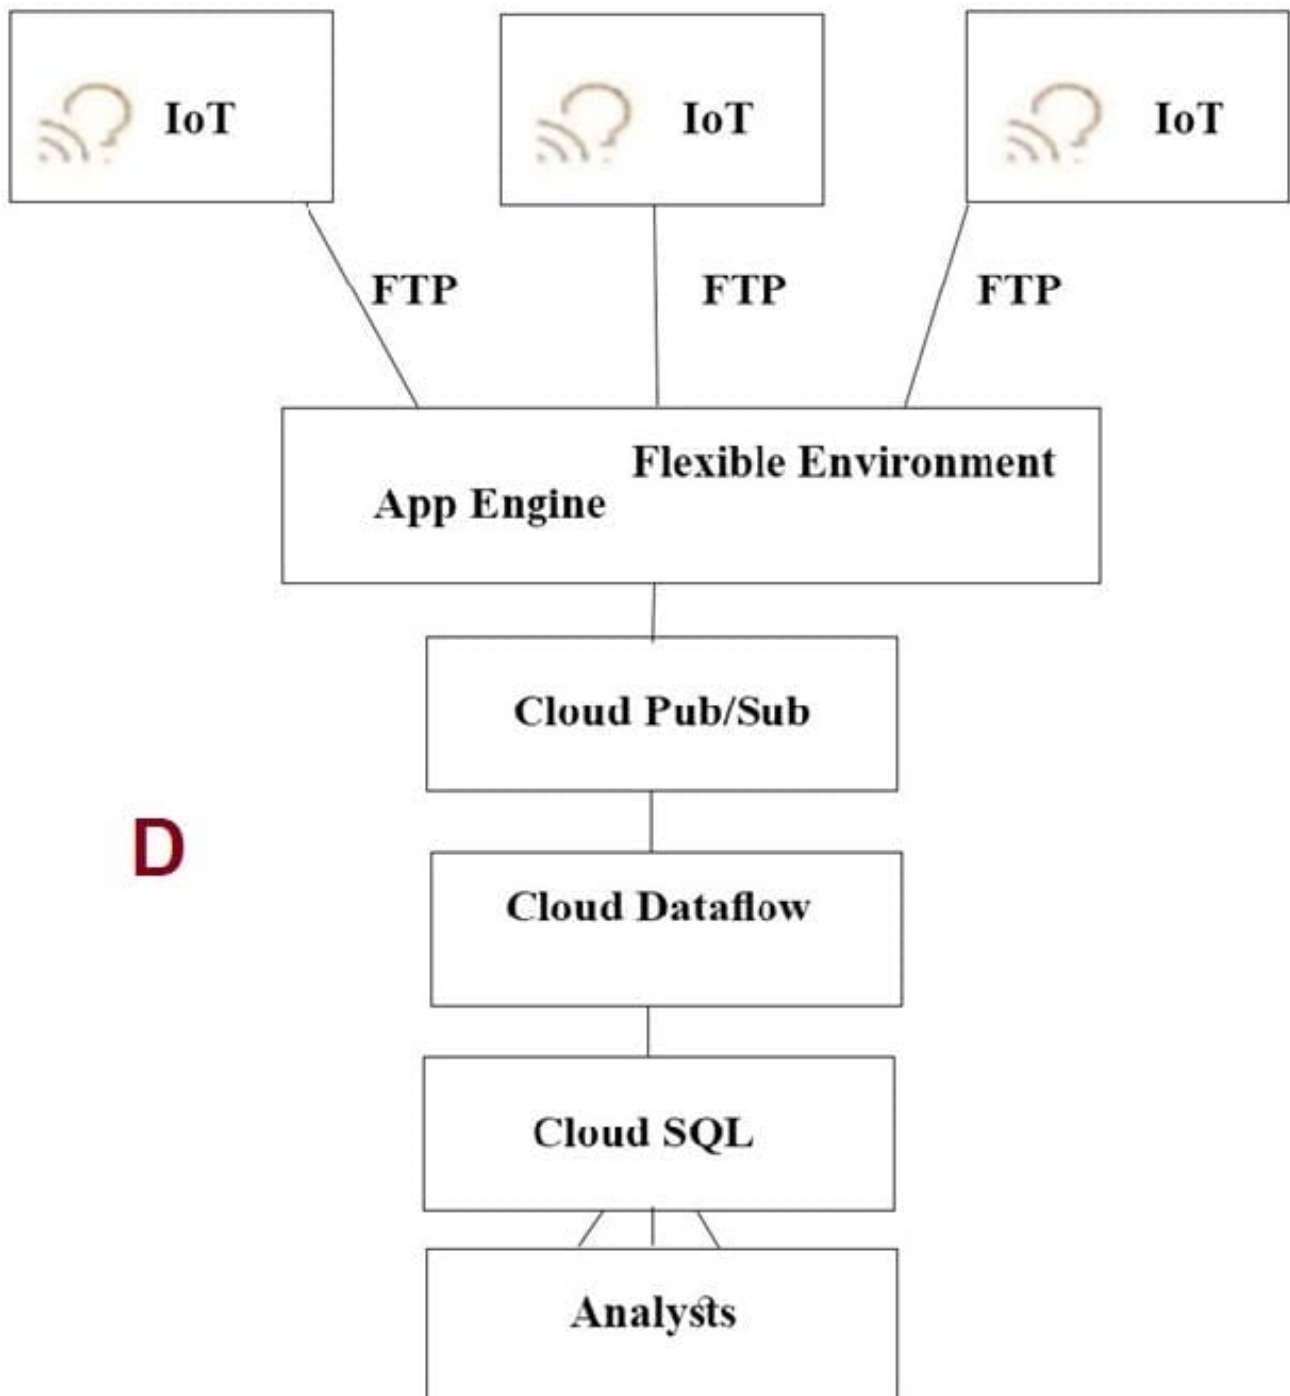
you might want to restrict external IP address so that only specific VM instances can use them. This option can help to prevent data exfiltration or maintain network isolation. Using an Organization Policy, you can restrict external IP addresses to specific VM instances with constraints to control use of external IP addresses for your VM instances within an organization or a project.

QUESTION 5

TerraEarth's CTO wants to use the raw data from connected vehicles to help identify approximately when a vehicle in the field will have a catastrophic failure. You want to allow analysts to centrally query the vehicle data.

Which architecture should you recommend?





A. Option A

B. Option B

C. Option C

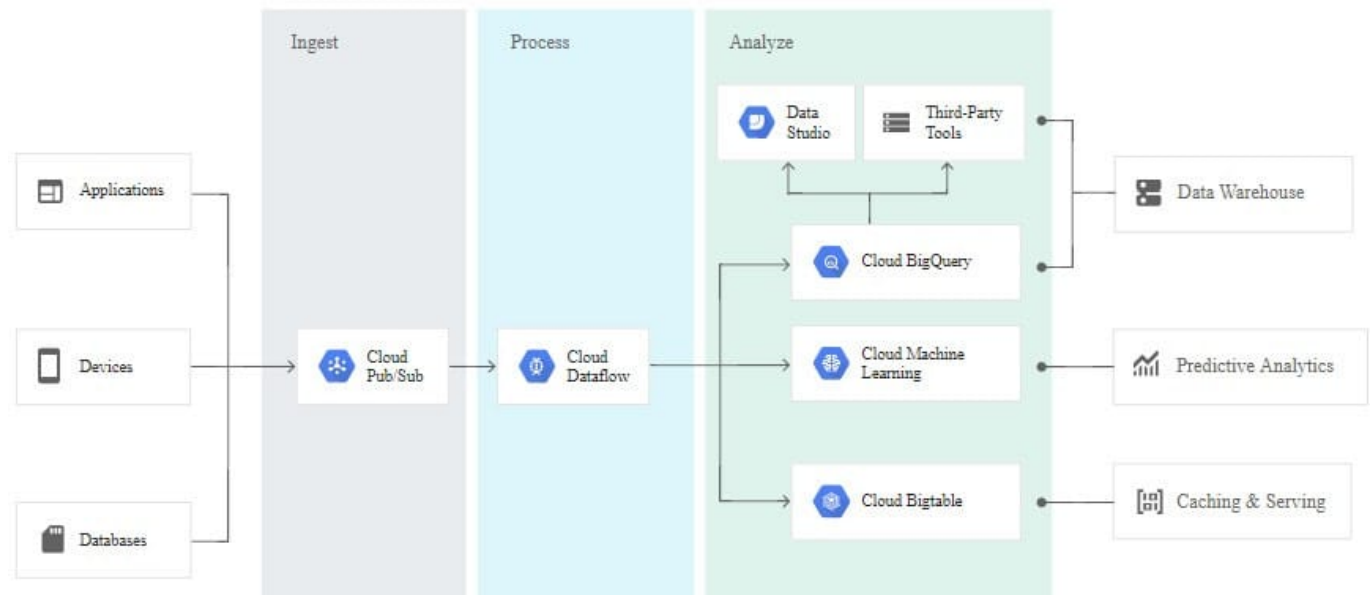
D. Option D

Correct Answer: A

The push endpoint can be a load balancer.

A container cluster can be used.

Cloud Pub/Sub for Stream Analytics



References:

<https://cloud.google.com/pubsub/>

<https://cloud.google.com/solutions/iot/>

<https://cloud.google.com/solutions/designing-connected-vehicle-platform>

https://cloud.google.com/solutions/designing-connected-vehicle-platform#data_ingestion

<http://www.eweek.com/big-data-and-analytics/google-touts-value-of-cloud-iot-core-for-analyzing-connected-car-data>

<https://cloud.google.com/solutions/iot/>

QUESTION 6

You are designing a large distributed application with 30 microservices. Each of your distributed microservices needs to connect to a database back-end. You want to store the credentials securely.

Where should you store the credentials?

- A. In the source code
- B. In an environment variable
- C. In a secret management system
- D. In a config file that has restricted access through ACLs



Correct Answer: C

References: <https://cloud.google.com/kms/docs/secret-management>

QUESTION 7

You need to migrate Hadoop jobs for your company's Data Science team without modifying the underlying infrastructure. You want to minimize costs and infrastructure management effort. What should you do?

- A. Create a Dataproc cluster using standard worker instances.
- B. Create a Dataproc cluster using preemptible worker instances.
- C. Manually deploy a Hadoop cluster on Compute Engine using standard instances.
- D. Manually deploy a Hadoop cluster on Compute Engine using preemptible instances.

Correct Answer: B

Reference: <https://cloud.google.com/architecture/hadoop/hadoop-gcp-migration-jobs>

QUESTION 8

You have an application that runs in Google Kubernetes Engine (GKE). Over the last 2 weeks, customers have reported that a specific part of the application returns errors very frequently. You currently have no logging or monitoring solution enabled on your GKE cluster. You want to diagnose the problem, but you have not been able to replicate the issue. You want to cause minimal disruption to the application. What should you do?

- A. 1. Update your GKE cluster to use Cloud Operations for GKE.
- 2. Use the GKE Monitoring dashboard to investigate logs from affected Pods.
- B. 1. Create a new GKE cluster with Cloud Operations for GKE enabled.
- 2.
- Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluster.
- 3.
- Use the GKE Monitoring dashboard to investigate logs from affected Pods.
- C. 1. Update your GKE cluster to use Cloud Operations for GKE, and deploy Prometheus.
- 2. Set an alert to trigger whenever the application returns an error.
- D. 1. Create a new GKE cluster with Cloud Operations for GKE enabled, and deploy Prometheus.
- 2.
- Migrate the affected Pods to the new cluster, and redirect traffic for those Pods to the new cluster.
- 3.



Set an alert to trigger whenever the application returns an error.

Correct Answer: A

Reference: <https://cloud.google.com/blog/products/management-tools/using-logging-your-apps-running-kubernetes-engine>

QUESTION 9

Your web application uses Google Kubernetes Engine to manage several workloads. One workload requires a consistent set of hostnames even after pod scaling and relaunches. Which feature of Kubernetes should you use to accomplish this?

- A. StatefulSets
- B. Role-based access control
- C. Container environment variables
- D. Persistent Volumes

Correct Answer: A

<https://kubernetes.io/docs/tutorials/stateful-application/basic-stateful-set/>

QUESTION 10

For this question, refer to the TerramEarth case study. You have broken down a legacy monolithic application into a few containerized RESTful microservices. You want to run those microservices on Cloud Run. You also want to make sure the services are highly available with low latency to your customers. What should you do?

- A. Deploy Cloud Run services to multiple availability zones. Create Cloud Endpoints that point to the services. Create a global HTTP(S) Load Balancing instance and attach the Cloud Endpoints to its backend.
- B. Deploy Cloud Run services to multiple regions. Create serverless network endpoint groups pointing to the services. Add the serverless NEG's to a backend service that is used by a global HTTP(S) Load Balancing instance.
- C. Deploy Cloud Run services to multiple regions. In Cloud DNS, create a latency-based DNS name that points to the services.
- D. Deploy Cloud Run services to multiple availability zones. Create a TCP/IP global load balancer. Add the Cloud Run Endpoints to its backend service.

Correct Answer: B

QUESTION 11

At Dress4Win, an operations engineer wants to create a low-cost solution to remotely archive copies of database backup files.



The database files are compressed tar files stored in their current data center.

How should he proceed?

- A. Create a cron script using gsutil to copy the files to a Coldline Storage bucket.
- B. Create a cron script using gsutil to copy the files to a Regional Storage bucket.
- C. Create a Cloud Storage Transfer Service Job to copy the files to a Coldline Storage bucket.
- D. Create a Cloud Storage Transfer Service job to copy the files to a Regional Storage bucket.

Correct Answer: C

As per the latest case study on google cloud website , they have DB storage of 1 PB out of which 600 TB is used. So you get the size of the data. These are the thumb rules as per GCP documentation Transferring from another cloud storage provider Use Storage Transfer Service

Transferring less than 1 TB from on-premises Use gsutil Transferring more than 1 TB from on-premises Use Transfer service for on-premises data <https://cloud.google.com/storage-transfer/docs/overview>

QUESTION 12

You have developed an application using Cloud ML Engine that recognizes famous paintings from uploaded images. You want to test the application and allow specific people to upload images for the next 24 hours. Not all users have a Google Account. How should you have users upload images?

- A. Have users upload the images to Cloud Storage. Protect the bucket with a password that expires after 24 hours.
- B. Have users upload the images to Cloud Storage using a signed URL that expires after 24 hours.
- C. Create an App Engine web application where users can upload images. Configure App Engine to disable the application after 24 hours. Authenticate users via Cloud Identity.
- D. Create an App Engine web application where users can upload images for the next 24 hours. Authenticate users via Cloud Identity.

Correct Answer: B

QUESTION 13

You have deployed an application on Anthos clusters (formerly Anthos GKE). According to the SRE practices at your company, you need to be alerted if request latency is above a certain threshold for a specified amount of time. What should you do?

- A. Install Anthos Service Mesh on your cluster. Use the Google Cloud Console to define a Service Level Objective (SLO), and create an alerting policy based on this SLO.
- B. Enable the Cloud Trace API on your project, and use Cloud Monitoring Alerts to send an alert based on the Cloud Trace metrics.
- C. Use Cloud Profiler to follow up the request latency. Create a custom metric in Cloud Monitoring based on the results of Cloud Profiler, and create an Alerting policy in case this metric exceeds the threshold.



D. Configure Anthos Config Management on your cluster, and create a yaml file that defines the SLO and alerting policy you want to deploy in your cluster.

Correct Answer: A

Reference: <https://cloud.google.com/anthos/docs/tutorials/manage-slos>

QUESTION 14

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 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.
- B. 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.
- C. 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.
- D. 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.

Correct Answer: C

QUESTION 15

Your company has successfully migrated to the cloud and wants to analyze their data stream to optimize operations. They do not have any existing code for this analysis, so they are exploring all their options. These options include a mix of batch and stream processing, as they are running some hourly jobs and live-processing some data as it comes in.

Which technology should they use for this?

- A. Google Cloud Dataproc
- B. Google Cloud Dataflow
- C. Google Container Engine with Bigtable
- D. Google Compute Engine with Google BigQuery

Correct Answer: B

Cloud Dataflow is a fully-managed service for transforming and enriching data in stream (real time) and batch (historical) modes with equal reliability and expressiveness -- no more complex workarounds or compromises needed. References: <https://cloud.google.com/dataflow/>



VCE & PDF

PassApply.com

<https://www.passapply.com/professional-cloud-architect.html>

2024 Latest passapply PROFESSIONAL-CLOUD-ARCHITECT PDF and VCE dumps Download

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

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

[PROFESSIONAL-CLOUD-ARCHITECT Braindumps](#)