



VCE & PDF

GeekCert.com

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

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

---

## 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



VCE & PDF

GeekCert.com

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

2024 Latest geekcert 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

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

Correct Answer: D

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/updating-apps> <https://cloud.google.com/blog/products/containers-kubernetes/kubernetes-best-practices-setting-up-health-checks-with-readiness-and-liveness-probes>

---

### QUESTION 2

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 instances.

Correct Answer: D

### QUESTION 3

You are working in a highly secured environment where public Internet access from the Compute Engine VMs is not



allowed. You do not yet have a VPN connection to access an on-premises file server. You need to install specific software on a Compute Engine instance. How should you install the software?

- A. Upload the required installation files to Cloud Storage. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gsutil.
- B. Upload the required installation files to Cloud Storage and use firewall rules to block all traffic except the IP address range for Cloud Storage. Download the files to the VM using gsutil.
- C. Upload the required installation files to Cloud Source Repositories. Configure the VM on a subnet with a Private Google Access subnet. Assign only an internal IP address to the VM. Download the installation files to the VM using gcloud.
- D. Upload the required installation files to Cloud Source Repositories and use firewall rules to block all traffic except the IP address range for Cloud Source Repositories. Download the files to the VM using gsutil.

Correct Answer: A

<https://cloud.google.com/vpc/docs/private-access-options#pga-supported>

#### QUESTION 4

For this question, refer to the EHR Healthcare case study. You need to define the technical architecture for hybrid connectivity between EHR's on-premises systems and Google Cloud. You want to follow Google's recommended practices for production-level applications. Considering the EHR Healthcare business and technical requirements, what should you do?

- A. Configure two Partner Interconnect connections in one metro (City), and make sure the Interconnect connections are placed in different metro zones.
- B. Configure two VPN connections from on-premises to Google Cloud, and make sure the VPN devices on-premises are in separate racks.
- C. Configure Direct Peering between EHR Healthcare and Google Cloud, and make sure you are peering at least two Google locations.
- D. Configure two Dedicated Interconnect connections in one metro (City) and two connections in another metro, and make sure the Interconnect connections are placed in different metro zones.

Correct Answer: D

based on the requirement of secure and high-performance connection between on-premises systems to Google Cloud  
<https://cloud.google.com/network-connectivity/docs/interconnect/tutorials/partner-creating-9999-availability>

#### QUESTION 5

Your company has an application running on multiple Compute Engine instances. You need to ensure that the application can communicate with an on-premises service that requires high throughput via internal IPs, while minimizing latency. What should you do?

- A. Use OpenVPN to configure a VPN tunnel between the on-premises environment and Google Cloud.
- B. Configure a direct peering connection between the on-premises environment and Google Cloud.



C. Use Cloud VPN to configure a VPN tunnel between the on-premises environment and Google Cloud.

D. Configure a Cloud Dedicated Interconnect connection between the on-premises environment and Google Cloud.

Correct Answer: D

Reference <https://cloud.google.com/architecture/setting-up-private-access-to-cloud-apis-through-vpn-tunnels>

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

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

[PROFESSIONAL-CLOUD-ARCHITECT Practice Test](#)