



**Google**

## **Exam Questions Professional-Cloud-Architect**

Google Certified Professional - Cloud Architect (GCP)

#### NEW QUESTION 1

- (Exam Topic 1)

For this question, refer to the Mountkirk Games case study.

Mountkirk Games wants you to design their new testing strategy. How should the test coverage differ from their existing backends on the other platforms?

- A. Tests should scale well beyond the prior approaches.
- B. Unit tests are no longer required, only end-to-end tests.
- C. Tests should be applied after the release is in the production environment.
- D. Tests should include directly testing the Google Cloud Platform (GCP) infrastructure.

**Answer:** A

#### Explanation:

From Scenario:

A few of their games were more popular than expected, and they had problems scaling their application servers, MySQL databases, and analytics tools.

Requirements for Game Analytics Platform include: Dynamically scale up or down based on game activity

#### NEW QUESTION 2

- (Exam Topic 1)

For this question, refer to the Mountkirk Games case study

Mountkirk Games needs to create a repeatable and configurable mechanism for deploying isolated application environments. Developers and testers can access each other's environments and resources, but they cannot access staging or production resources. The staging environment needs access to some services from production.

What should you do to isolate development environments from staging and production?

- A. Create a project for development and test and another for staging and production.
- B. Create a network for development and test and another for staging and production.
- C. Create one subnetwork for development and another for staging and production.
- D. Create one project for development, a second for staging and a third for production.

**Answer:** D

#### NEW QUESTION 3

- (Exam Topic 1)

For this question, refer to the Mountkirk Games case study.

Mountkirk Games wants to set up a continuous delivery pipeline. Their architecture includes many small services that they want to be able to update and roll back quickly. Mountkirk Games has the following requirements:

- Services are deployed redundantly across multiple regions in the US and Europe.
- Only frontend services are exposed on the public internet.
- They can provide a single frontend IP for their fleet of services.
- Deployment artifacts are immutable. Which set of products should they use?

- A. Google Cloud Storage, Google Cloud Dataflow, Google Compute Engine
- B. Google Cloud Storage, Google App Engine, Google Network Load Balancer
- C. Google Kubernetes Registry, Google Container Engine, Google HTTP(S) Load Balancer
- D. Google Cloud Functions, Google Cloud Pub/Sub, Google Cloud Deployment Manager

**Answer:** C

#### NEW QUESTION 4

- (Exam Topic 5)

You are helping the QA team to roll out a new load-testing tool to test the scalability of your primary cloud services that run on Google Compute Engine with Cloud Bigtable. Which three requirements should they include? Choose 3 answers

- A. Ensure that the load tests validate the performance of Cloud Bigtable.
- B. Create a separate Google Cloud project to use for the load-testing environment.
- C. Schedule the load-testing tool to regularly run against the production environment.
- D. Ensure all third-party systems your services use are capable of handling high load.
- E. Instrument the production services to record every transaction for replay by the load-testing tool.
- F. Instrument the load-testing tool and the target services with detailed logging and metrics collection.

**Answer:** ABF

#### NEW QUESTION 5

- (Exam Topic 5)

You have been asked to select the storage system for the click-data of your company's large portfolio of websites. This data is streamed in from a custom website analytics package at a typical rate of 6,000 clicks per minute, with bursts of up to 8,500 clicks per second. It must be stored for future analysis by your data science and user experience teams. Which storage infrastructure should you choose?

- A. Google Cloud SQL
- B. Google Cloud Bigtable
- C. Google Cloud Storage
- D. Google cloud Datastore

**Answer:** C

#### Explanation:

<https://cloud.google.com/bigquery/docs/loading-data-cloud-storage>

#### NEW QUESTION 6

- (Exam Topic 5)

Your company's test suite is a custom C++ application that runs tests throughout each day on Linux virtual machines. The full test suite takes several hours to complete, running on a limited number of on premises servers reserved for testing. Your company wants to move the testing infrastructure to the cloud, to reduce the amount of time it takes to fully test a change to the system, while changing the tests as little as possible. Which cloud infrastructure should you recommend?

- A. Google Compute Engine unmanaged instance groups and Network Load Balancer
- B. Google Compute Engine managed instance groups with auto-scaling
- C. Google Cloud Dataproc to run Apache Hadoop jobs to process each test
- D. Google App Engine with Google Stackdriver for logging

**Answer:** B

#### Explanation:

<https://cloud.google.com/compute/docs/instance-groups/>

Google Compute Engine enables users to launch virtual machines (VMs) on demand. VMs can be launched from the standard images or custom images created by users.

Managed instance groups offer autoscaling capabilities that allow you to automatically add or remove instances from a managed instance group based on increases or decreases in load. Autoscaling helps your applications gracefully handle increases in traffic and reduces cost when the need for resources is lower.

#### NEW QUESTION 7

- (Exam Topic 5)

You want to establish a Compute Engine application in a single VPC across two regions. The application must communicate over VPN to an on-premises network. How should you deploy the VPN?

- A. Use VPC Network Peering between the VPC and the on-premises network.
- B. Expose the VPC to the on-premises network using IAM and VPC Sharing.
- C. Create a global Cloud VPN Gateway with VPN tunnels from each region to the on-premises peer gateway.
- D. Deploy Cloud VPN Gateway in each regio
- E. Ensure that each region has at least one VPN tunnel to the on-premises peer gateway.

**Answer:** C

#### Explanation:

<https://cloud.google.com/vpn/docs/how-to/creating-static-vpns>

#### NEW QUESTION 8

- (Exam Topic 5)

You are tasked with building an online analytical processing (OLAP) marketing analytics and reporting tool. This requires a relational database that can operate on hundreds of terabytes of data. What is the Google recommended tool for such applications?

- A. Cloud Spanner, because it is globally distributed
- B. Cloud SQL, because it is a fully managed relational database
- C. Cloud Firestore, because it offers real-time synchronization across devices
- D. BigQuery, because it is designed for large-scale processing of tabular data

**Answer:** A

#### Explanation:

Reference: <https://cloud.google.com/files/BigQueryTechnicalWP.pdf>

#### NEW QUESTION 9

- (Exam Topic 5)

A development manager is building a new application He asks you to review his requirements and identify what cloud technologies he can use to meet them. The application must

- \* 1. Be based on open-source technology for cloud portability
- \* 2. Dynamically scale compute capacity based on demand
- \* 3. Support continuous software delivery
- \* 4. Run multiple segregated copies of the same application stack
- \* 5. Deploy application bundles using dynamic templates
- \* 6. Route network traffic to specific services based on URL

Which combination of technologies will meet all of his requirements?

- A. Google Container Engine, Jenkins, and Helm
- B. Google Container Engine and Cloud Load Balancing
- C. Google Compute Engine and Cloud Deployment Manager
- D. Google Compute Engine, Jenkins, and Cloud Load Balancing

**Answer:** A

#### Explanation:

Helm for managing Kubernetes

Kubernetes can base on the URL to route traffic to different location (path) <https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer> eg.apiVersion: networking.k8s.io/v1beta1

kind: Ingress metadata:

name: fanout-ingress spec:

rules:

- http: paths:
- path: /\* backend: serviceName: web servicePort: 8080
- path: /v2/\* backend: serviceName: web2 servicePort: 8080

#### NEW QUESTION 10

- (Exam Topic 5)

Your web application has several VM instances running within a VPC. You want to restrict communications between instances to only the paths and ports you authorize, but you don't want to rely on static IP addresses or subnets because the app can autoscale. How should you restrict communications?

- A. Use separate VPCs to restrict traffic
- B. Use firewall rules based on network tags attached to the compute instances
- C. Use Cloud DNS and only allow connections from authorized hostnames
- D. Use service accounts and configure the web application particular service accounts to have access

**Answer:** B

#### NEW QUESTION 10

- (Exam Topic 5)

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 Folder for each department
- D. A single Organization with multiple projects, each with a central owner

**Answer:** C

#### Explanation:

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>

#### NEW QUESTION 15

- (Exam Topic 5)

You are analyzing and defining business processes to support your startup's trial usage of GCP, and you don't yet know what consumer demand for your product will be. Your manager requires you to minimize GCP service costs and adhere to Google best practices. What should you do?

- A. Utilize free tier and sustained use discount
- B. Provision a staff position for service cost management.
- C. Utilize free tier and sustained use discount
- D. Provide training to the team about service cost management.
- E. Utilize free tier and committed use discount
- F. Provision a staff position for service cost management.
- G. Utilize free tier and committed use discount
- H. Provide training to the team about service cost management.

**Answer:** D

#### Explanation:

[https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#billing\\_and\\_management](https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#billing_and_management)

#### NEW QUESTION 18

- (Exam Topic 5)

A production database virtual machine on Google Compute Engine has an ext4-formatted persistent disk for data files The database is about to run out of storage space How can you remediate the problem with the least amount of downtime?

- A. In the Cloud Platform Console, increase the size of the persistent disk and use the `resize2fs` command in Linux.
- B. Shut down the virtual machine, use the Cloud Platform Console to increase the persistent disk size, then restart the virtual machine.
- C. In the Cloud Platform Console, increase the size of the persistent disk and verify the new space is ready to use with the `fdisk` command in Linux.
- D. In the Cloud Platform Console, create a new persistent disk attached to the virtual machine, format and mount it, and configure the database service to move the files to the new disk.
- E. In the Cloud Platform Console, create a snapshot of the persistent disk, restore the snapshot to a new larger disk, unmount the old disk, mount the new disk, and restart the database service.

**Answer:** A

#### Explanation:

On Linux instances, connect to your instance and manually resize your partitions and file systems to use the additional disk space that you added.

Extend the file system on the disk or the partition to use the added space. If you grew a partition on your disk, specify the partition. If your disk does not have a partition table, specify only the disk ID.

```
sudo resize2fs /dev/[DISK_ID][PARTITION_NUMBER]
```

where [DISK\_ID] is the device name and [PARTITION\_NUMBER] is the partition number for the device where you are resizing the file system.

References: <https://cloud.google.com/compute/docs/disks/add-persistent-disk>

#### NEW QUESTION 23

- (Exam Topic 5)

You need to upload files from your on-premises environment to Cloud Storage. You want the files to be encrypted on Cloud Storage using customer-supplied encryption keys. What should you do?

- A. Supply the encryption key in a .boto configuration file
- B. Use gsutil to upload the files.
- C. Supply the encryption key using gcloud config
- D. Use gsutil to upload the files to that bucket.
- E. Use gsutil to upload the files, and use the flag --encryption-key to supply the encryption key.
- F. Use gsutil to create a bucket, and use the flag --encryption-key to supply the encryption key
- G. Use gsutil to upload the files to that bucket.

**Answer:** A

**Explanation:**

<https://cloud.google.com/storage/docs/encryption/customer-supplied-keys#gsutil>

#### NEW QUESTION 27

- (Exam Topic 5)

Google Cloud Platform resources are managed hierarchically using organization, folders, and projects. When Cloud Identity and Access Management (IAM) policies exist at these different levels, what is the effective policy at a particular node of the hierarchy?

- A. The effective policy is determined only by the policy set at the node
- B. The effective policy is the policy set at the node and restricted by the policies of its ancestors
- C. The effective policy is the union of the policy set at the node and policies inherited from its ancestors
- D. The effective policy is the intersection of the policy set at the node and policies inherited from its ancestors

**Answer:** B

**Explanation:**

Reference: <https://cloud.google.com/resource-manager/docs/cloud-platform-resource-hierarchy>

#### NEW QUESTION 29

- (Exam Topic 5)

You have been engaged by your client to lead the migration of their application infrastructure to GCP. One of their current problems is that the on-premises high performance SAN is requiring frequent and expensive upgrades to keep up with the variety of workloads that are identified as follows: 20TB of log archives retained for legal reasons; 500 GB of VM boot/data volumes and templates; 500 GB of image thumbnails; 200 GB of customer session state data that allows customers to restart sessions even if off-line for several days.

Which of the following best reflects your recommendations for a cost-effective storage allocation?

- A. Local SSD for customer session state data
- B. Lifecycle-managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.
- C. Memcache backed by Cloud Datastore for the customer session state data
- D. Lifecycle-managed Cloud Storage for log archives, thumbnails, and VM boot/data volumes.
- E. Memcache backed by Cloud SQL for customer session state data
- F. Assorted local SSD-backed instances for VM boot/data volume
- G. Cloud Storage for log archives and thumbnails.
- H. Memcache backed by Persistent Disk SSD storage for customer session state data
- I. Assorted local SSD-backed instances for VM boot/data volume
- J. Cloud Storage for log archives and thumbnails.

**Answer:** D

**Explanation:**

<https://cloud.google.com/compute/docs/disks>

#### NEW QUESTION 30

- (Exam Topic 5)

Your company has decided to build a backup replica of their on-premises user authentication PostgreSQL database on Google Cloud Platform. The database is 4 TB, and large updates are frequent. Replication requires private address space communication. Which networking approach should you use?

- A. Google Cloud Dedicated Interconnect
- B. Google Cloud VPN connected to the data center network
- C. A NAT and TLS translation gateway installed on-premises
- D. A Google Compute Engine instance with a VPN server installed connected to the data center network

**Answer:** A

**Explanation:**

<https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations>

Google Cloud Dedicated Interconnect provides direct physical connections and RFC 1918 communication between your on-premises network and Google's network. Dedicated Interconnect enables you to transfer large amounts of data between networks, which can be more cost effective than purchasing additional bandwidth over the public Internet or using VPN tunnels.

Benefits:

➤ Traffic between your on-premises network and your VPC network doesn't traverse the public Internet.

Traffic traverses a dedicated connection with fewer hops, meaning there are less points of failure where traffic might get dropped or disrupted.

➤ Your VPC network's internal (RFC 1918) IP addresses are directly accessible from your on-premises network. You don't need to use a NAT device or VPN tunnel to reach internal IP addresses. Currently, you can only reach internal IP addresses over a dedicated connection. To reach Google external IP addresses, you must use a separate connection.

- You can scale your connection to Google based on your needs. Connection capacity is delivered over one or more 10 Gbps Ethernet connections, with a maximum of eight connections (80 Gbps total per interconnect).
- The cost of egress traffic from your VPC network to your on-premises network is reduced. A dedicated connection is generally the least expensive method if you have a high-volume of traffic to and from Google's network.
- References: <https://cloud.google.com/interconnect/docs/details/dedicated>

#### NEW QUESTION 34

- (Exam Topic 5)

You have a Python web application with many dependencies that requires 0.1 CPU cores and 128 MB of memory to operate in production. You want to monitor and maximize machine utilization. You also to reliably deploy new versions of the application. Which set of steps should you take?

- A. Perform the following:1) Create a managed instance group with f1-micro type machines.2) Use a startup script to clone the repository, check out the production branch, install the dependencies, and start the Python app.3) Restart the instances to automatically deploy new production releases.
- B. Perform the following:1) Create a managed instance group with n1-standard-1 type machines.2) Build a Compute Engine image from the production branch that contains all of the dependencies and automatically starts the Python app.3) Rebuild the Compute Engine image, and update the instance template to deploy new production releases.
- C. Perform the following:1) Create a Kubernetes Engine cluster with n1-standard-1 type machines.2) Build a Docker image from the production branch with all of the dependencies, and tag it with the version number.3) Create a Kubernetes Deployment with the imagePullPolicy set to "IfNotPresent" in the staging namespace, and then promote it to the production namespace after testing.
- D. Perform the following:1) Create a Kubernetes Engine cluster with n1-standard-4 type machines.2) Build a Docker image from the master branch will all of the dependencies, and tag it with "latest".3) Create a Kubernetes Deployment in the default namespace with the imagePullPolicy set to "Always". Restart the pods to automatically deploy new production releases.

**Answer: D**

#### Explanation:

<https://cloud.google.com/compute/docs/instance-templates>

#### NEW QUESTION 38

- (Exam Topic 5)

You want to enable your running Google Kubernetes Engine cluster to scale as demand for your application changes. What should you do?

- A. Add additional nodes to your Kubernetes Engine cluster using the following command:`gcloud container clusters resizeCLUSTER_Name --size 10`
- B. Add a tag to the instances in the cluster with the following command:`gcloud compute instances add-tagsINSTANCE - -tags enable-autoscaling max-nodes-10`
- C. Update the existing Kubernetes Engine cluster with the following command:`gcloud alpha container clustersupdate mycluster - -enable-autoscaling - -min-nodes=1 - -max-nodes=10`
- D. Create a new Kubernetes Engine cluster with the following command:`gcloud alpha container clusterscreate mycluster - -enable-autoscaling - -min-nodes=1 - -max-nodes=10`and redeploy your application

**Answer: C**

#### Explanation:

<https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler> To enable autoscaling for an existing node pool, run the following command:  
`gcloud container clusters update [CLUSTER_NAME] --enable-autoscaling --min-nodes 1 --max-nodes 10 --zone [COMPUTE_ZONE] --node-pool default-pool`

#### NEW QUESTION 39

- (Exam Topic 5)

You want to automate the creation of a managed instance group and a startup script to install the OS package dependencies. You want to minimize the startup time for VMs in the instance group. What should you do?

- A. Use Terraform to create the managed instance group and a startup script to install the OS package dependencies.
- B. Create a custom VM image with all OS package dependencies
- C. Use Deployment Manager to create the managed instance group with the VM image.
- D. Use Puppet to create the managed instance group and install the OS package dependencies.
- E. Use Deployment Manager to create the managed instance group and Ansible to install the OS package dependencies.

**Answer: B**

#### Explanation:

"Custom images are more deterministic and start more quickly than instances with startup scripts. However, startup scripts are more flexible and let you update the apps and settings in your instances more easily." [https://cloud.google.com/compute/docs/instance-templates/create-instance-templates#using\\_custom\\_or\\_public\\_i](https://cloud.google.com/compute/docs/instance-templates/create-instance-templates#using_custom_or_public_i)

#### NEW QUESTION 40

- (Exam Topic 5)

You are deploying a PHP App Engine Standard service with SQL as the backend. You want to minimize the number of queries to the database. What should you do?

- A. Set the memcache service level to dedicate
- B. Create a key from the hash of the query, and return database values from memcache before issuing a query to Cloud SQL.
- C. Set the memcache service level to dedicate
- D. Create a cron task that runs every minute to populate the cache with keys containing query results.
- E. Set the memcache service level to share
- F. Create a cron task that runs every minute to save all expected queries to a key called "cached-queries".
- G. Set the memcache service level to share
- H. Create a key called "cached-queries", and return database values from the key before using a query to Cloud SQL.

**Answer:** A

**Explanation:**

<https://cloud.google.com/appengine/docs/standard/php/memcache/using>

#### NEW QUESTION 42

- (Exam Topic 5)

Your customer support tool logs all email and chat conversations to Cloud Bigtable for retention and analysis. What is the recommended approach for sanitizing this data of personally identifiable information or payment card information before initial storage?

- A. Hash all data using SHA256
- B. Encrypt all data using elliptic curve cryptography
- C. De-identify the data with the Cloud Data Loss Prevention API
- D. Use regular expressions to find and redact phone numbers, email addresses, and credit card numbers

**Answer:** A

**Explanation:**

Reference: <https://cloud.google.com/solutions/pci-dss-compliance-ingcp#>

#### NEW QUESTION 46

- (Exam Topic 5)

Your company is running a stateless application on a Compute Engine instance. The application is used heavily during regular business hours and lightly outside of business hours. Users are reporting that the application is slow during peak hours. You need to optimize the application's performance. What should you do?

- A. Create a snapshot of the existing dis
- B. Create an instance template from the snapsho
- C. Create an autoscaled managed instance group from the instance template.
- D. Create a snapshot of the existing dis
- E. Create a custom image from the snapsho
- F. Create an autoscaled managed instance group from the custom image.
- G. Create a custom image from the existing dis
- H. Create an instance template from the custom image.Create an autoscaled managed instance group from the instance template.
- I. Create an instance template from the existing dis
- J. Create a custom image from the instance template.Create an autoscaled managed instance group from the custom image.

**Answer:** B

**Explanation:**

<https://cloud.google.com/compute/docs/instance-templates/create-instance-templates>

#### NEW QUESTION 48

- (Exam Topic 5)

Your development team has installed a new Linux kernel module on the batch servers in Google Compute Engine (GCE) virtual machines (VMs) to speed up the nightly batch process. Two days after the installation, 50% of web application deployed in the same nightly batch run. You want to collect details on the failure to pass back to the development team. Which three actions should you take? Choose 3 answers

- A. Use Stackdriver Logging to search for the module log entries.
- B. Read the debug GCE Activity log using the API or Cloud Console.
- C. Use gcloud or Cloud Console to connect to the serial console and observe the logs.
- D. Identify whether a live migration event of the failed server occurred, using in the activity log.
- E. Adjust the Google Stackdriver timeline to match the failure time, and observe the batch server metrics.
- F. Export a debug VM into an image, and run the image on a local server where kernel log messages will be displayed on the native screen.

**Answer:** ACE

**Explanation:**

<https://www.flexera.com/blog/cloud/2013/12/google-compute-engine-live-migration-passes-the-test/> "With live migration, the virtual machines are moved without any downtime or noticeable service degradation"

#### NEW QUESTION 51

- (Exam Topic 5)

Your company wants to track whether someone is present in a meeting room reserved for a scheduled meeting. There are 1000 meeting rooms across 5 offices on 3 continents. Each room is equipped with a motion sensor that reports its status every second. The data from the motion detector includes only a sensor ID and several different discrete items of information. Analysts will use this data, together with information about account owners and office locations. Which database type should you use?

- A. Flat file
- B. NoSQL
- C. Relational
- D. Blobstore

**Answer:** B

**Explanation:**

Relational databases were not designed to cope with the scale and agility challenges that face modern applications, nor were they built to take advantage of the commodity storage and processing power available today.

NoSQL fits well for:

Developers are working with applications that create massive volumes of new, rapidly changing data types — structured, semi-structured, unstructured and polymorphic data.

#### NEW QUESTION 56

- (Exam Topic 5)

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 archived files and files uploaded daily using the `gsutil -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
- C. Establish a connection with Google using a Dedicated Interconnect or Direct Peering connection and use it to upload files daily.
- D. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transfer archived data to Cloud Storage
- E. Establish one Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily using the `gsutil -m` option.
- F. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storage
- G. Establish a Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily.

**Answer:** B

**Explanation:**

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

#### NEW QUESTION 61

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### Professional-Cloud-Architect Practice Exam Features:

- \* Professional-Cloud-Architect Questions and Answers Updated Frequently
- \* Professional-Cloud-Architect Practice Questions Verified by Expert Senior Certified Staff
- \* Professional-Cloud-Architect Most Realistic Questions that Guarantee you a Pass on Your First Try
- \* Professional-Cloud-Architect Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The Professional-Cloud-Architect Practice Test Here](#)**