

Valid Professional-Cloud-Architect Exam Dumps Ensure you a HIGH SCORE (2024) [Q63-Q86]



Valid Professional-Cloud-Architect Exam Dumps Ensure you a HIGH SCORE (2024) Pass Professional-Cloud-Architect Exam with Latest Questions

Google Professional-Cloud-Architect certification exam has been designed to validate the skills of cloud professionals who are interested in designing and developing solutions on Google Cloud Platform. Candidates who pass the exam will have demonstrated their expertise in designing and implementing cloud solutions that are scalable, reliable, and secure. Professional-Cloud-Architect exam covers a wide range of topics, including cloud architecture, infrastructure, security, and compliance. Candidates are required to demonstrate their knowledge of these topics by answering multiple-choice questions, case studies, and scenarios.

Google Professional-Cloud-Architect certification is an excellent way for professionals to enhance their career prospects in the cloud computing industry. It demonstrates to potential employers that the individual has the skills and knowledge required to design, develop, and manage cloud-based solutions on the Google Cloud Platform. It also helps professionals stand out from their peers and demonstrates their commitment to learning and staying up-to-date with the latest cloud computing technologies.

NO.63 For this question, refer to the Helicopter Racing League (HRL) case study. A recent finance audit of cloud infrastructure noted an exceptionally high number of Compute Engine instances are allocated to do video encoding and transcoding. You suspect that these Virtual Machines are zombie machines that were not deleted after their workloads completed. You need to quickly get a list of which VM instances are idle. What should you do?

- * Log into each Compute Engine instance and collect disk, CPU, memory, and network usage statistics for analysis.
- * Use the `gcloud compute instances list` to list the virtual machine instances that have the `idle: true` label set.
- * Use the `gcloud recommender` command to list the idle virtual machine instances.
- * From the Google Console, identify which Compute Engine instances in the managed instance groups are no longer responding to

health check probes.

Reference: <https://cloud.google.com/compute/docs/instances/viewing-and-applying-idle-vm-recommendations>

NO.64 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?

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

<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>

NO.65 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?

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

NO.66 You need to ensure reliability for your application and operations by supporting reliable task scheduling for compute on GCP. Leveraging Google best practices, what should you do?

- * Using the Cron service provided by App Engine, publishing messages directly to a message-processing utility service running on Compute Engine instances.
- * Using the Cron service provided by App Engine, publish messages to a Cloud Pub/Sub topic.

Subscribe to that topic using a message-processing utility service running on Compute Engine instances.

- * Using the Cron service provided by Google Kubernetes Engine (GKE), publish messages directly to a message-processing utility service running on Compute Engine instances.
- * Using the Cron service provided by GKE, publish messages to a Cloud Pub/Sub topic. Subscribe to that topic using a message-processing utility service running on Compute Engine instances.

NO.67 Your company wants to start using Google Cloud resources but wants to retain their on-premises Active Directory domain controller for identity management. What should you do?

- * Use the Admin Directory API to authenticate against the Active Directory domain controller.
- * Use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML SSO.
- * Use Cloud Identity-Aware Proxy configured to use the on-premises Active Directory domain controller as an identity provider.
- * Use Compute Engine to create an Active Directory (AD) domain controller that is a replica of the on-premises AD domain controller using Google Cloud Directory Sync.

Explanation

https://cloud.google.com/solutions/federating-gcp-with-active-directory-introduction#implementing_federation

NO.68 Your organization has a 3-tier web application deployed in the same network on Google Cloud Platform.

Each tier (web, API, and database) scales independently of the others. Network traffic should flow through the web to the API tier and then on to the database tier. Traffic should not flow between the web and the database tier.

How should you configure the network?

- * Add each tier to a different subnetwork
- * Set up software based firewalls on individual VMs
- * Add tags to each tier and set up routes to allow the desired traffic flow
- * Add tags to each tier and set up firewall rules to allow the desired traffic flow

Google Cloud Platform(GCP) enforces firewall rules through rules and tags. GCP rules and tags can be defined once and used across all regions.

Reference: <https://cloud.google.com/docs/compare/openstack/>

<https://aws.amazon.com/it/blogs/aws/building-three-tier-architectures-with-security-groups/>

NO.69 Case Study: 5 – Dress4win

Company Overview

Dress4win is a web-based company that helps their users organize and manage their personal wardrobe using a website and mobile application. The company also cultivates an active social network that connects their users with designers and retailers. They monetize their services through advertising, e-commerce, referrals, and a freemium app model. The application has grown from a few servers in the founder’s garage to several hundred servers and appliances in a collocated data center. However, the capacity of their infrastructure is now insufficient for the application’s rapid growth. Because of this growth and the company’s desire to innovate faster.

Dress4Win is committing to a full migration to a public cloud.

Solution Concept

For the first phase of their migration to the cloud, Dress4win is moving their development and test environments. They are also building a disaster recovery site, because their current infrastructure is at a single location. They are not sure which components of

their architecture they can migrate as is and which components they need to change before migrating them.

Existing Technical Environment

The Dress4win application is served out of a single data center location. All servers run Ubuntu LTS v16.04.

Databases:

MySQL. 1 server for user data, inventory, static data:

- * MySQL 5.8

- 8 core CPUs

- 128 GB of RAM

- 2x 5 TB HDD (RAID 1)

Redis 3 server cluster for metadata, social graph, caching. Each server is:

- * Redis 3.2

- 4 core CPUs

- 32GB of RAM

Compute:

40 Web Application servers providing micro-services based APIs and static content.

- * Tomcat
- Java

- Nginx

- 4 core CPUs

- 32 GB of RAM

20 Apache Hadoop/Spark servers:

- * Data analysis

- Real-time trending calculations

- 8 core CPUs

- 128 GB of RAM

- 4x 5 TB HDD (RAID 1)

3 RabbitMQ servers for messaging, social notifications, and events:

- * 8 core CPUs

- * 32GB of RAM

Miscellaneous servers:

- * Jenkins, monitoring, bastion hosts, security scanners

- * 8 core CPUs

- * 32GB of RAM

Storage appliances:

iSCSI for VM hosts

- * Fiber channel SAN; MySQL databases

- * 1 PB total storage; 400 TB available

NAS; image storage, logs, backups

- * 100 TB total storage; 35 TB available

Business Requirements

Build a reliable and reproducible environment with scaled parity of production.

- * Improve security by defining and adhering to a set of security and Identity and Access

- * Management (IAM) best practices for cloud.

Improve business agility and speed of innovation through rapid provisioning of new resources.

- * Analyze and optimize architecture for performance in the cloud.

- * Technical Requirements

Easily create non-production environment in the cloud.

- * Implement an automation framework for provisioning resources in cloud.

- * Implement a continuous deployment process for deploying applications to the on-premises

- * datacenter or cloud.

Support failover of the production environment to cloud during an emergency.

- * Encrypt data on the wire and at rest.
- * Support multiple private connections between the production data center and cloud environment.

Executive Statement

Our investors are concerned about our ability to scale and contain costs with our current infrastructure. They are also concerned that a competitor could use a public cloud platform to offset their up-front investment and free them to focus on developing better features. Our traffic patterns are highest in the mornings and weekend evenings; during other times, 80% of our capacity is sitting idle.

Our capital expenditure is now exceeding our quarterly projections. Migrating to the cloud will likely cause an initial increase in spending, but we expect to fully transition before our next hardware refresh cycle. Our total cost of ownership (TCO) analysis over the next 5 years for a public cloud strategy achieves a cost reduction between 30% and 50% over our current model.

For this question, refer to the Dress4Win case study. Dress4Win is expected to grow to 10 times its size in 1 year with a corresponding growth in data and traffic that mirrors the existing patterns of usage. The CIO has set the target of migrating production infrastructure to the cloud within the next 6 months. How will you configure the solution to scale for this growth without making major application changes and still maximize the ROI?

- * Migrate the web application layer to App Engine, and MySQL to Cloud Datastore, and NAS to Cloud Storage. Deploy RabbitMQ, and deploy Hadoop servers using Deployment Manager.
- * Migrate RabbitMQ to Cloud Pub/Sub, Hadoop to BigQuery, and NAS to Compute Engine with Persistent Disk storage. Deploy Tomcat, and deploy Nginx using Deployment Manager.
- * Implement managed instance groups for Tomcat and Nginx. Migrate MySQL to Cloud SQL, RabbitMQ to Cloud Pub/Sub, Hadoop to Cloud Dataproc, and NAS to Compute Engine with Persistent Disk storage.
- * Implement managed instance groups for the Tomcat and Nginx. Migrate MySQL to Cloud SQL, RabbitMQ to Cloud Pub/Sub, Hadoop to Cloud Dataproc, and NAS to Cloud Storage.

NO.70 For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to design their solution for the future in order to take advantage of cloud and technology improvements as they become available. Which two steps should they take? (Choose two.)

- * Store as much analytics and game activity data as financially feasible today so it can be used to train machine learning models to predict user behavior in the future.
- * Begin packaging their game backend artifacts in container images and running them on Google Kubernetes Engine to improve the availability to scale up or down based on game activity.
- * Set up a CI/CD pipeline using Jenkins and Spinnaker to automate canary deployments and improve development velocity.
- * Adopt a schema versioning tool to reduce downtime when adding new game features that require storing additional player data in the database.
- * Implement a weekly rolling maintenance process for the Linux virtual machines so they can apply critical kernel patches and package updates and reduce the risk of 0-day vulnerabilities.

NO.71 Operational parameters such as oil pressure are adjustable on each of TerramEarth's vehicles to increase their efficiency, depending on their environmental conditions. Your primary goal is to increase the operating efficiency of all 20 million cellular and unconnected vehicles in the field.

How can you accomplish this goal?

- * Have your engineers inspect the data for patterns, and then create an algorithm with rules that make operational adjustments automatically

- * Capture all operating data, train machine learning models that identify ideal operations, and run locally to make operational adjustments automatically
- * Implement a Google Cloud Dataflow streaming job with a sliding window, and use Google Cloud Messaging (GCM) to make operational adjustments automatically
- * Capture all operating data, train machine learning models that identify ideal operations, and host in Google Cloud Machine Learning (ML) Platform to make operational adjustments automatically

Explanation/Reference: <https://cloud.google.com/customers/ocado/>

TerramEarth, B

Testlet 1

Company Overview

TerramEarth manufactures heavy equipment for the mining and agricultural industries. About 80% of their business is from mining and 20% from agriculture. They currently have over 500 dealers and service centers in

100 countries. Their mission is to build products that make their customers more productive.

Solution Concept

There are 20 million TerramEarth vehicles in operation that collect 120 fields of data per second. Data is stored locally on the vehicle and can be accessed for analysis when a vehicle is serviced. The data is downloaded via a maintenance port. This same port can be used to adjust operational parameters, allowing the vehicles to be upgraded in the field with new computing modules.

Approximately 200,000 vehicles are connected to a cellular network, allowing TerramEarth to collect data directly. At a rate of 120 fields of data per second, with 22 hours of operation per day, TerramEarth collects a total of about 9 TB/day from these connected vehicles.

Existing Technical Environment

TerramEarth's existing architecture is composed of Linux and Windows-based systems that reside in a single

U.S. west coast based data center. These systems gzip CSV files from the field and upload via FTP, and place the data in their data warehouse. Because this process takes time, aggregated reports are based on data that is 3 weeks old.

With this data, TerramEarth has been able to preemptively stock replacement parts and reduce unplanned downtime of their vehicles by 60%. However, because the data is stale, some customers are without their vehicles for up to 4 weeks while they wait for replacement parts.

Business Requirements

- * Decrease unplanned vehicle downtime to less than 1 week
- * Support the dealer network with more data on how their customers use their equipment to better position new products and services
- * Have the ability to partner with different companies; especially with seed and fertilizer suppliers in the fast-growing agricultural business; to create compelling joint offerings for their customers

- * Expand beyond a single datacenter to decrease latency to the American Midwest and east coast
- * Create a backup strategy
- * Increase security of data transfer from equipment to the datacenter
- * Improve data in the data warehouse
- * Use customer and equipment data to anticipate customer needs

Application 1: Data ingest

A custom Python application reads uploaded datafiles from a single server, writes to the data warehouse.

Compute:

- * Windows Server 2008 R2

– 16 CPUs

– 128 GB of RAM

– 10 TB local HDD storage

Application 2: Reporting

An off the shelf application that business analysts use to run a daily report to see what equipment needs repair.

Only 2 analysts of a team of 10 (5 west coast, 5 east coast) can connect to the reporting application at a time.

Compute:

- * Off the shelf application. License tied to number of physical CPUs

– Windows Server 2008 R2

– 16 CPUs

– 32 GB of RAM

– 500 GB HDD

Data warehouse:

- * A single PostgreSQL server

– RedHat Linux

– 64 CPUs

• 128 GB of RAM

• 4x 6TB HDD in RAID 0

Executive Statement

Our competitive advantage has always been in our manufacturing process, with our ability to build better vehicles for lower cost than our competitors. However, new products with different approaches are constantly being developed, and I'm concerned that we lack the skills to undergo the next wave of transformations in our industry. My goals are to build our skills while addressing immediate market needs through incremental innovations.

NO.72 One of your primary business objectives is being able to trust the data stored in your application. You want to log all changes to the application data.

How can you design your logging system to verify authenticity of your logs?

- * Write the log concurrently in the cloud and on premises
- * Use a SQL database and limit who can modify the log table
- * Digitally sign each timestamp and log entry and store the signature
- * Create a JSON dump of each log entry and store it in Google Cloud Storage

Write a log entry. If the log does not exist, it is created. You can specify a severity for the log entry, and you can write a structured log entry by specifying `–payload-type=json` and writing your message as a JSON string:

```
gcloud logging write LOG STRING
```

```
gcloud logging write LOG JSON-STRING &#8211;payload-type=json
```

Reference: <https://cloud.google.com/logging/docs/reference/tools/gcloud-logging>

NO.73 You have created several preemptible Linux virtual machine instances using Google Compute Engine.

You want to properly shut down your application before the virtual machines are preempted. What should you do?

- * Create a shutdown script named `k99.shutdown` in the `/etc/rc.6.d/` directory.
- * Create a shutdown script registered as a `xinetd` service in Linux and configure a `Stackdriver` endpoint check to call the service.
- * Create a shutdown script and use it as the value for a new metadata entry with the key `shutdown-script` in the Cloud Platform Console when you create the new virtual machine instance.
- * Create a shutdown script, registered as a `xinetd` service in Linux, and use the `gcloud compute instances add-metadata` command to specify the service URL as the value for a new metadata entry with the key `shutdown-script-url`

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

NO.74 Your solution is producing performance bugs in production that you did not see in staging and test environments. You want to adjust your test and deployment procedures to avoid this problem in the future. What should you do?

- * Deploy fewer changes to production.
- * Deploy smaller changes to production.
- * Increase the load on your test and staging environments.
- * Deploy changes to a small subset of users before rolling out to production.

NO.75 You are using Cloud Shell and need to install a custom utility for use in a few weeks. Where can you store the file so it is in the default execution path and persists across sessions?

- * `~/bin`
- * Cloud Storage

- * /google/scripts
- * /usr/local/bin

Explanation

NO.76 For this question, refer to the TerramEarth case study

You analyzed TerramEarth's business requirement to reduce downtime, and found that they can achieve a majority of time saving by reducing customers' wait time for parts. You decided to focus on reduction of the 3 weeks aggregate reporting time. Which modifications to the company's processes should you recommend?

- * Migrate from CSV to binary format, migrate from FTP to SFTP transport, and develop machine learning analysis of metrics.
- * Migrate from FTP to streaming transport, migrate from CSV to binary format, and develop machine learning analysis of metrics.
- * Increase fleet cellular connectivity to 80%, migrate from FTP to streaming transport, and develop machine learning analysis of metrics.
- * Migrate from FTP to SFTP transport, develop machine learning analysis of metrics, and increase dealer local inventory by a fixed factor.

Reference:

The Avro binary format is the preferred format for loading compressed data. Avro data is faster to load because the data can be read in parallel, even when the data blocks are compressed.

Cloud Storage supports streaming transfers with the gsutil tool or boto library, based on HTTP chunked transfer encoding. Streaming data lets you stream data to and from your Cloud Storage account as soon as it becomes available without requiring that the data be first saved to a separate file. Streaming transfers are useful if you have a process that generates data and you do not want to buffer it locally before uploading it, or if you want to send the result from a computational pipeline directly into Cloud Storage.

References: <https://cloud.google.com/storage/docs/streaming>

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

NO.77 You are using Cloud Shell and need to install a custom utility for use in a few weeks. Where can you store the file so it is in the default execution path and persists across sessions?

- * ~/bin
- * Cloud Storage
- * /google/scripts
- * /usr/local/bin

NO.78 For this question, refer to the Mountkirk Games case study. You are in charge of the new Game Backend Platform architecture. The game communicates with the backend over a REST API.

You want to follow Google-recommended practices. How should you design the backend?

- * Create an instance template for the backend. For every region, deploy it on a multi-zone managed instance group. Use an L4 load balancer.
- * Create an instance template for the backend. For every region, deploy it on a single-zone managed instance group. Use an L4 load balancer.
- * Create an instance template for the backend. For every region, deploy it on a multi-zone managed instance group. Use an L7 load balancer.
- * Create an instance template for the backend. For every region, deploy it on a single-zone managed instance group. Use an L7 load balancer.

Explanation

https://cloud.google.com/solutions/gaming/cloud-game-infrastructure#dedicated_game_server

NO.79 An application development team believes their current logging tool will not meet their needs for their new cloud-based product. They want a better tool to capture errors and help them analyze their historical log data. You want to help them find a solution that meets their needs, what should you do?

- * Direct them to download and install the Google StackDriver logging agent.
- * Send them a list of online resources about logging best practices.
- * Help them define their requirements and assess viable logging tools.
- * Help them upgrade their current tool to take advantage of any new features.

Help them define their requirements and assess viable logging tools. They know the requirements and the existing tools' problems. While its true StackDriver Logging and Error Reporting possibly meet all their requirements, there might be other tools also meet their need. They need you to provide expertise to make assessment for new tools, specifically, logging tools that can capture errors and help them analyze their historical log data;

References:

<https://cloud.google.com/logging/docs/agent/installation>

NO.80 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?

- * Google Cloud SQL
- * Google Cloud Bigtable
- * Google Cloud Storage
- * Google cloud Datastore

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

NO.81 Case Study: 2 ; TerramEarth Case Study

Company Overview

TerramEarth manufactures heavy equipment for the mining and agricultural industries: About 80% of their business is from mining and 20% from agriculture. They currently have over 500 dealers and service centers in 100 countries. Their mission is to build products that make their customers more productive.

Company Background

TerramEarth formed in 1946, when several small, family owned companies combined to retool after World War II. The company cares about their employees and customers and considers them to be extended members of their family.

TerramEarth is proud of their ability to innovate on their core products and find new markets as their customers' needs change. For the past 20 years trends in the industry have been largely toward increasing productivity by using larger vehicles with a human operator.

Solution Concept

There are 20 million TerramEarth vehicles in operation that collect 120 fields of data per second.

Data is stored locally on the vehicle and can be accessed for analysis when a vehicle is serviced.

The data is downloaded via a maintenance port. This same port can be used to adjust operational parameters, allowing the vehicles to be upgraded in the field with new computing modules.

Approximately 200,000 vehicles are connected to a cellular network, allowing TerramEarth to collect data directly. At a rate of 120 fields of data per second, with 22 hours of operation per day.

TerramEarth collects a total of about 9 TB/day from these connected vehicles.

Existing Technical Environment



TerramEarth's existing architecture is composed of Linux-based systems that reside in a data center. These systems gzip CSV files from the field and upload via FTP, transform and aggregate them, and place the data in their data warehouse. Because this process takes time, aggregated reports are based on data that is 3 weeks old.

With this data, TerramEarth has been able to preemptively stock replacement parts and reduce unplanned downtime of their vehicles by 60%. However, because the data is stale, some customers are without their vehicles for up to 4 weeks while they wait for replacement parts.

Business Requirements

• Decrease unplanned vehicle downtime to less than 1 week, without

increasing the cost of carrying surplus inventory

• Support the dealer network with more data on how their customers use

their equipment IP better position new products and services.

• Have the ability to partner with different companies-especially with

seed and fertilizer suppliers in the fast-growing agricultural

business-to create compelling joint offerings for their customers

CEO Statement

We have been successful in capitalizing on the trend toward larger vehicles to increase the productivity of our customers.

Technological change is occurring rapidly and TerramEarth has taken advantage of connected devices technology to provide our

customers with better services, such as our intelligent farming equipment. With this technology, we have been able to increase farmers' yields by 25%, by using past trends to adjust how our vehicles operate. These advances have led to the rapid growth of our agricultural product line, which we expect will generate 50% of our revenues by 2020.

CTO Statement

Our competitive advantage has always been in the manufacturing process with our ability to build better vehicles for lower cost than our competitors. However, new products with different approaches are constantly being developed, and I'm concerned that we lack the skills to undergo the next wave of transformations in our industry. Unfortunately, our CEO doesn't take technology obsolescence seriously and he considers the many new companies in our industry to be niche players. My goals are to build our skills while addressing immediate market needs through incremental innovations.

For this question, refer to the TerraEarth case study. TerraEarth has equipped unconnected trucks with servers and sensors to collect telemetry data

- * Next year they want to use the data to train machine learning models. They want to store this data in the cloud while reducing costs. What should they do?
- * Have the vehicle's computer compress the data in hourly snapshots, and store it in a Google Cloud storage (GCS) Nearline bucket.
- * Push the telemetry data in Real-time to a streaming dataflow job that compresses the data, and store it in Google BigQuery.
- * Push the telemetry data in real-time to a streaming dataflow job that compresses the data, and store it in Cloud Bigtable.
- * Have the vehicle's computer compress the data in hourly snapshots, and store it in a GCS Coldline bucket.

Storage is the best choice for data that you plan to access at most once a year, due to its slightly lower availability, 90-day minimum storage duration, costs for data access, and higher per-operation costs. For example:

Cold Data Storage; Infrequently accessed data, such as data stored for legal or regulatory reasons, can be stored at low cost as Coldline Storage, and be available when you need it.

Disaster recovery; In the event of a disaster recovery event, recovery time is key. Cloud Storage provides low latency access to data stored as Coldline Storage.

References: <https://cloud.google.com/storage/docs/storage-classes>

NO.82 Your organization has a 3-tier web application deployed in the same network on Google Cloud Platform. Each tier (web, API, and database) scales independently of the others. Network traffic should flow through the web to the API tier and then on to the database tier.

Traffic should not flow between the web and the database tier. How should you configure the network?

- * Add each tier to a different subnet.
- * Set up software based firewalls on individual VMs.
- * Add tags to each tier and set up routes to allow the desired traffic flow.
- * Add tags to each tier and set up firewall rules to allow the desired traffic flow.

NO.83 To reduce costs, the Director of Engineering has required all developers to move their development infrastructure resources from on-premises virtual machines (VMs) to Google Cloud Platform. These resources go through multiple start/stop events during the day and require state to persist. You have been asked to design the process of running a development environment in Google Cloud while providing cost visibility to the finance department. Which two steps should you take? Choose 2 answers

- * Use the `--no-auto-delete` flag on all persistent disks and stop the VM.
- * Use the `-auto-delete` flag on all persistent disks and terminate the VM.
- * Apply VM CPU utilization label and include it in the BigQuery billing export.
- * Use Google BigQuery billing export and labels to associate cost to groups.

- * Store all state into local SSD, snapshot the persistent disks, and terminate the VM.
- * Store all state in Google Cloud Storage, snapshot the persistent disks, and terminate the VM.

Explanation

<https://cloud.google.com/billing/docs/how-to/export-data-bigquery>

Team or cost center labels: Add labels based on team or cost center to distinguish instances owned by different teams (for example, team:research and team:analytics). You can use this type of label for cost accounting or budgeting.

<https://cloud.google.com/resource-manager/docs/creating-managing-labels>

NO.84 Dress4Win would like to become familiar with deploying applications to the cloud by successfully deploying some applications quickly, as is. They have asked for your recommendation.

What should you advise?

- * Identify self-contained applications with external dependencies as a first move to the cloud.
- * Identify enterprise applications with internal dependencies and recommend these as a first move to the cloud.
- * Suggest moving their in-house databases to the cloud and continue serving requests to on-premise applications.
- * Recommend moving their message queuing servers to the cloud and continue handling requests to on-premise applications.

NO.85 A few days after JencoMart migrates the user credentials database to Google Cloud Platform and shuts down the old server, the new database server stops responding to SSH connections. It is still serving database requests to the application servers correctly.

What three steps should you take to diagnose the problem? Choose 3 answers.

- * Delete the virtual machine (VM) and disks and create a new one
- * Delete the instance, attach the disk to a new VM, and investigate
- * Take a snapshot of the disk and connect to a new machine to investigate
- * Check inbound firewall rules for the network the machine is connected to
- * Connect the machine to another network with very simple firewall rules and investigate
- * Print the Serial Console output for the instance for troubleshooting, activate the interactive console, and investigate

D: Handling “Unable to connect on port 22” error message

Possible causes include:

- * There is no firewall rule allowing SSH access on the port. SSH access on port 22 is enabled on all Compute Engine instances by default. If you have disabled access, SSH from the Browser will not work. If you run sshd on a port other than 22, you need to enable the access to that port with a custom firewall rule.
- * The firewall rule allowing SSH access is enabled, but is not configured to allow connections from GCP Console services. Source IP addresses for browser-based SSH sessions are dynamically allocated by GCP Console and can vary from session to session.

F: Handling “Could not connect, retrying…” error

You can verify that the daemon is running by navigating to the serial console output page and looking for output lines prefixed with the accounts-from-metadata: string. If you are using a standard image but you do not see these output prefixes in the serial console output, the daemon might be stopped. Reboot the instance to restart the daemon.

References:

<https://cloud.google.com/compute/docs/ssh-in-browser>

<https://cloud.google.com/compute/docs/ssh-in-browser>

NO.86 For this question, refer to the JencoMart case study.

JencoMart has built a version of their application on Google Cloud Platform that serves traffic to Asia. You want to measure success against their business and technical goals. Which metrics should you track?

- * Error rates for requests from Asia
- * Latency difference between US and Asia
- * Total visits, error rates, and latency from Asia
- * Total visits and average latency for users in Asia
- * The number of character sets present in the database

Reference:

The benefit of obtaining the Google Professional Cloud Architect Exam Certification - Google Professional Cloud Architect certification has more useful and relevant networks that help them in setting career goals for themselves. Google Professional Cloud Architect networks provide them with the correct career guidance than non-certified generally are unable to get.- Google Professional Cloud Architect Certification is distinguished among competitors. Google Professional Cloud Architect certification can give them an edge at that time easily when candidates appear for employment interview, employers are very fascinated to note one thing that differentiates the individual from all other candidates.**Google Professional Cloud Architect Exam** provide proven knowledge to use the tools to complete the task efficiently and cost effectively than the other non-certified professionals lack in doing so.- Google Professional Cloud Architect will be confident and stand different from others as their skills are more trained than non-certified professionals. **Professional-Cloud-Architect Exam**

Practice Questions prepared by Google Professionals:

https://www.braindumpsit.com/Professional-Cloud-Architect_real-exam.html