Answer Sheet

Q1)

An app for a finance company needs access to a database and a Cloud Storage bucket.

There is no predefined role that grants all the needed permissions without granting some permissions that are not needed.

You decide to create a custom role.

When defining custom roles, you should follow which of the following principles?

- Rotation of duties
- Hierarchical inheritance
- Defense in depth
- ✅ Least privilege

Explanation:-The principle of least privilege states that users should have only the privileges that are needed to carry out their duties. Rotation of duties means that different people should perform a task at different times. Defense in depth is the practice of using multiple security controls to protect the same asset. is not a real security principal. Hierarchical inheritance means that policies at the organization level, the folder level, the project level, or at the resource level are inherited by all its child resources. For any further detail: https://cloud.google.com/iam/docs/using-iam-securely

---

Q2)

You work in an international Company based in North America and your boss told you that you have to plan the GDPR compliance for EUROPE.

Which of the following elements you have to take care of:

- ✅ Create an updated inventory of personal data that you handle.

Explanation:-The General Data Protection Regulation (EU) 2016/679 (GDPR) is a regulation in EU law on data protection and privacy in the European Union (EU) and the European Economic Area (EEA). It also addresses the transfer of personal data outside the EU and EEA areas. The GDPR aims primarily to give control to individuals over their personal data and to simplify the regulatory environment for international business by unifying the regulation within the EU. B is wrong because GDPR protects privacy, so it is against of unsolicited distribution of private information B is wrong because CCPA are different For any further detail: https://en.wikipedia.org/wiki/General_Data_Protection_Regulation https://cloud.google.com/security/gdpr/

- Let the common data to be public, if the Customer doesn't advice against
- Use California Consumer Protection Act (CCPA) rules
- ✅ Review your current controls, policies, and processes for managing and protecting data

Explanation:-The General Data Protection Regulation (EU) 2016/679 (GDPR) is a regulation in EU law on data protection and privacy in the European Union (EU) and the European Economic Area (EEA). It also addresses the transfer of personal data outside the EU and EEA areas. The GDPR aims primarily to give control to individuals over their personal data and to simplify the regulatory environment for international business by unifying the regulation within the EU. B is wrong because GDPR protects privacy, so it is against of unsolicited distribution of private information B is wrong because CCPA are different For any further detail: https://en.wikipedia.org/wiki/General_Data_Protection_Regulation https://cloud.google.com/security/gdpr/

Q3)

Your company is planning to migrate to GCP. Migration of SQL databases needs to be addressed.

Your company is international and therefore needs a global database for integrated ERP and local databases for specific applications for each country.

All data must be integrable and will be stored into an Analytics System, too.

All services must be fully managed.

The on-premises Systems used to have primary keys represented by a sequence, a progressive number.

What kind of row key do you choose?

- ✅ Combination of fields not starting with a progressive number

Explanation:-There are 2 issues: you have to choose keys on a common ground; that is better for integration of different systems. So you have to meet Spanner requirements; Spanner is a SQL DB but is global, splitted architecture; you have to avoid hotspots, that is, splits that are crowded and others that are hardly accessed. Cloud Spanner assigns work to different servers in units of splits, so the server assigned to this particular split ends up handling all the insert requests. As the frequency of user access events increases, the frequency of insert requests to the corresponding server also increases. The server then becomes prone to becoming a hotspot. So you cannot use close values or numbers. A good choice is Version 4 UUID, too, because it uses random values in the bit sequence. Version 1 UUID stores the timestamp in the high order bits and is not recommended. For any further detail: https://cloud.google.com/spanner/docs/schema-design

- String Fields only
- Progressive numbers
- Timestamp of inserting time
- ✅ Universally Unique Identifier ver 4

Explanation:-There are 2 issues: you have to choose keys on a common ground; that is better for integration of different systems. So you have to meet Spanner requirements; Spanner is a SQL DB but is global, splitted architecture; you have to avoid hotspots, that is, splits that are crowded and others that are hardly accessed. Cloud Spanner assigns work to different servers in units of splits, so the server assigned to this particular split ends up handling all the insert requests. As the frequency of user access events increases, the frequency of insert requests to the corresponding server also increases. The server then becomes prone to becoming a hotspot. So you cannot use close values or numbers. A good choice is Version 4 UUID, too, because it uses random values in the bit sequence. Version 1 UUID stores the timestamp in the high order bits and is not recommended. For any further detail: https://cloud.google.com/spanner/docs/schema-design

---

Q4)

A team of mobile developers is developing a new application.

It will require synchronizing data between mobile devices and a backend database.

Which database service would you recommend?

- Cloud SQL
- BigQuery
- ✅ Firestore

Explanation:-Firestore, part of GCP and of Firebase is the only Database designed for Web and Mobile Application that provides live synchronization and offline support, Cloud Firestore is a fast, fully managed, serverless, cloud-native NoSQL document database that simplifies storing, syncing, and querying data for mobile, web, and IoT apps at global scale. Cloud Firestore is the next generation of Cloud Datastore. So Datastore is just the same of Firestore. For any further detail: https://cloud.google.com/firestore/

- Spanner
- Bigtable

Q5)

You are migrating a series of applications to Google Cloud Platform with a lift and shift methodology, using Compute Engine.

Applications must be scalable, so Load Balancer and instance groups are being configured.

Some applications manage session data in memory.

Which of the following configurations do you choose to allow apps to work properly:

- ✅ HTTP(S) load balancer with Session affinity

Explanation:-Session affinity provides a best-effort attempt to send requests from a particular client to the same backend for as long as the backend is healthy and has the capacity, according to the configured balancing mode. It is the best way to assure that session data is maintained in memory. It is a feature of the HTTP(S) load balancing. Google Cloud SSL Proxy Load Balancing terminates user SSL (TLS) connections at the load balancing layer, then balances the connections across your instances using the SSL or TCP protocols. Cloud SSL proxy is intended for non-HTTP(S) traffic. For HTTP(S) traffic, HTTP(S) load balancing is recommended instead. pass-through load balancer, so your backends receive the original client request For any further detail: https://cloud.google.com/load-balancing/docs/https/ https://cloud.google.com/load-balancing/docs/https/#session_affinity https://cloud.google.com/load-balancing/docs/ssl/ https://cloud.google.com/load-balancing/docs/network/

- HTTP(S) load balancer with WebSocket proxy support
- QUIC protocol support for HTTPS Load Balancing
- Network Load Balancing
- SSL Proxy con Health Checks

Q6)

You have been asked to migrate to the Cloud, in a “lift and shift way”, a legacy system, You have set up a managed Instance Groups and Load Balancers.

You need to set up health checks: which kind of resources are addressed and which are the parameters that you have to insert (pick 1)?

- Load Balancer, name, check frequency
- Load Balancer, name, protocol and port
- VMs, name, check frequency
- ✅ VMs, name, protocol and port

Explanation:-Health checks monitor VMs when used with load balancers; you need to provide an health check name and protocol and port to be checked. The health check works in the following way: The metadata server at IP address 169.254.169.254 is responsible for sending traffic to the health check URL. The destination address of the health check is the load balancer's external address. This process mimics real incoming traffic. The health check must be answered with an HTTP 200 status followed by a normal TCP connection closure within the time specified by the timeoutSec setting. For any further detail: https://cloud.google.com/load-balancing/docs/health-check-concepts

---

Q7)

Your team is developing a microservice backend that provides an Application Programming interface with security and control functionality.

In order to organize test development and production in a better way, the team asked you to find a managed, simple and safe solution that could speed up the time in compliance with the requirements.

Note that the systems are developed with different techniques and operate in GKE clusters.

- Apigee Hybrid
- ✅ Cloud Endpoints for OpenAPI

Explanation:-Cloud Endpoints is a GCP API management system. It provides an API console, hosting, logging, monitoring, and other features to create, share, maintain, and secure your APIs. Cloud Endpoints uses the distributed Extensible Service Proxy (ESP) to provide low latency and high performance for serving APIs. ESP is a service proxy based on NGINX, that runs in its own Docker container for better isolation and scalability and is distributed in the Container Registry. It can be used with App Engine flexible, Google Kubernetes Engine (GKE), Compute Engine or Kubernetes. A is out of scope because it is not a Hybrid solution Apigee hybrid can host API traffic on-premises, Google Cloud, or a hybrid with an enterprise-grade hybrid gateway. C is wrong because we are not using the App Engine standard environment D is wrong because Cloud Tasks lets separate out pieces of work that can be performed independently with handlers (tasks). For any further detail: https://cloud.google.com/endpoints/docs/openapi/about-cloud-endpoints https://cloud.google.com/endpoints/docs/openapi/get-started-kubernetes-engineCloud Endpoints Frameworks for the App Engine standard environment

- Cloud Endpoints for the App Engine standard environment
- Cloud Tasks

Q8)

Your company is planning to migrate to the Cloud a large part of its System.

You have to make a presentation of the main GCP Security mechanisms that are not currently used in your company that can better protect services and resources.

Which of the following solutions are you going to present (pick 3)?

- GCP uses only automatic security reviews ; you may activate your own manual security reviews , fuzzers, static analysis tools, and web security scanners for XSS and other vulnerabilities
- Google uses firewalling and network segmentation as primary security mechanisms
- ✅ Each GCP Service has an associated service account identity with cryptographic credentials

Explanation:-Each service that runs on the infrastructure has an associated service account identity. A service is provided cryptographic credentials that it can use to prove its identity when making or receiving remote procedure calls (RPCs) to other services. These identities are used by clients to ensure that they are talking to the correct intended server, and by servers to limit access to methods and data to particular clients.

- All external communications are under cryptographic protection. Optionally you may use hardware cryptographic accelerators for traffic inside GCP data centers
- ✅ Services open to the Internet can be registered to Google Front End (GFE), that ensures TLS certificates and overall secrecy

Explanation:-When a service wants to make itself available on the Internet, it can register itself with an infrastructure service called the Google Front End (GFE). The GFE ensures that all TLS connections are terminated using correct certificates and following best practices such as supporting perfect forward secrecy. The GFE additionally applies protections against Denial of Service attacks (which we will discuss in more detail later). The GFE then forwards requests for the service using the RPC security protocol discussed previously. B is false because GCP We not rely on internal network segmentation or firewalling as our primary security mechanisms, though it do use ingress and egress filtering at various points in our network to prevent IP spoofing as a further security layer. This approach also helps to maximize our network's performance and availability. D is false: Services can configure the level of cryptographic protection they want for each infrastructure RPC G is false: Google Cloud Platform encrypts customer data stored at rest by default, with no additional action required from you F is wrong Beyond the central source control GCP also provides manual security reviews that range from quick triages for less risky features to in-depth design and implementation reviews for the most risky features. These reviews are conducted by a team that includes experts across web security, cryptography, and operating system security. The reviews can also result in new security library features and new fuzzers that can then be applied to other future products. For any further detail: https://cloud.google.com/security/infrastructure/design/

- ✅ Google use cryptographic signatures over low-level components (from BIOS to OS images) to ensure security on the correct software used and each Physical Server has its own identity

Explanation:-From documentation: A: Google server machines use a variety of technologies to ensure that they are booting the correct software stack. We use cryptographic signatures over low-level components like the BIOS, bootloader, kernel, and base operating system image. These signatures can be validated during each boot or update. The components are all Google-controlled, built, and hardened. With each new generation of hardware we strive to continually improve security: for example, depending on the generation of server design, we root the trust of the boot chain in either a lockable firmware chip, a microcontroller running Google-written security code, or the above mentioned Google-designed security chip.

Q9) Which are the purpose and differences between the following Security Modules? HSM, KMS and Secret Manager

- HSM manages the complete lifecycle of a key KMS make just the storage Secret Manager talks with the applications
- ✅ HSM is a hardware security module that encrypt and store security keys, KMS manage and rotate cryptographic keys for your cloud services and Secret Manager manage passwords, credentials, API keys for your cloud services

Explanation:-Google's infrastructure provides a variety of storage services, such as Bigtable and Spanner, and a central key management service: Cloud KMS. Most applications at Google access physical storage indirectly via these storage services. The storage services can be configured to use keys from the central key management service to encrypt data before it is written to physical storage. This key management service supports automatic key rotation, provides extensive audit logs, and integrates with the previously mentioned end user permission tickets to link keys to particular end users. So, Google Cloud Key Management Service (KMS) is a cloud service for managing automatically all the services related to encryption keys for other Google cloud services that enterprises can use to implement cryptographic functions. HSM is a physical computing device that stores and manages digital keys for strong authentication and provides crypto-processing. They usually plug-in cards or external devices that are attached directly to a computer or network server. Cloud HSM is a managed service for HSM and it is fully integrated with KMS for creating and using customer-managed encryption keys. It is necessary only in special cases where an hardware enforced additional level of security is required. Secrets are database credentials, passwords, keys secrets, any security token. rotate, manage, and retrieve secrets. Secrets Manager helps you protect secrets needed to access your applications, services, and IT resources. It is needed when you store secrets but manage yourself the security procedures. For any further detail: https://cloud.google.com/hsm/ https://cloud.google.com/kms/ https://cloud.google.com/secret-manager/docs/

- HSM manage passwords, credentials, API keys for your cloud services KMS is a hardware security module that encrypts and store security keys Secret Manager manage and rotate cryptographic keys for your cloud services and Secret Manager manage passwords
- HSM manage and rotate cryptographic keys for your cloud services and Secret Manager manage passwords, KMS manage passwords, credentials, API keys for your cloud services, Secret Manager ia an hardware security module that encrypt and store security keys and

---

Q10)

You have just completed a Python app. At this point you need to: Deploy it on App Engine Standard Expose it with a custom domain name Activate SSL/TLS

What do you have to do:

- Nothing, App Engine provides everything you need
- You cannot use a custom domain name but only http: // [YOUR_PROJECT_ID] .appspot.com
- You must purchase both a domain name and an https certificate
- ✅ You must purchase a domain name but not an https certificate

Explanation:-App Engine has a full SSL support; it fives globally-distributed SSL endpoints and built-in load balancing to serve your app securely, reliably, and quickly to a worldwide audience. By default, HTTPS connections on your custom domain will be enabled automatically using managed SSL certificates. Once a custom domain is mapped to an application and DNS records are configured, App Engine provisions a managed SSL certificate, handles renewing the certificate before it expires, and revokes it when you remove the custom domain from your application. The other answers are so, clearly wrong. For any further detail: https://cloud.google.com/appengine/docs/standard/python/securing-custom-domains-with-ssl https://cloud.google.com/appengine/docs/standard/python/getting-started/deploying-the-application

---

Q11) You have an app that runs in a VM in GCP Compute Engine.
You have been asked to organize the development / test process to make Automating builds from Source Code.
Which product should you use and with what configuration?

- Cloud RUN with Dockerfile and a build trigger
- ✅ Cloud Build with Dockerfile or a Cloud Build build config file and a build trigger

Explanation:-The Correct Product is Cloud Build, that is a service that executes your builds on Google Cloud Platform infrastructure. Cloud Build can import source code from Google Cloud Storage, Cloud Source Repositories, GitHub, or Bitbucket, execute a build to your specifications, and produce artifacts such as Docker containers or Java archives.
App Engine and GKE may use Cloud Build to realize a continuous integration and continuous deployment solution, so the other answers are wrong without the Cloud Build product and triggers.
For any further detail: https://cloud.google.com/cloud-build/docs/running-builds/automate-builds

- GKE with Spinnaker
- App Engine with Cloud Source Repositories

Q12)

A client of yours has asked you for advice because he is looking for a quick and convenient solution for adding functionalities to an Application.

Whenever a new customer is created in the Firebase database, he wants to perform a series of welcome activities and a series of follow-up actions, regardless of the specific function that recorded the new customer record.

Which of the following solutions will you suggest (choose 1)?

- Compute Engine and Managed Instances
- App Engine Flexible Environment
- App Engine Standard Environment
- ✅ Cloud Functions

Explanation:-Any of these environments can host the additional functionalities required. The best solution is with Cloud Function, because you can handle events in the Firebase Realtime Database with no need to update client code. Cloud Functions may have the full administrative privileges to ensures that each change to the database is processed individually. Furthermore Cloud Functions are a decoupled economic solution, because of the pay-as-you go model. Functions handle database events in 2 ways; listening for specifically for only creation, update, or deletion events, or you can listen for any change of any kind to a path. These Cloud Functions event handlers are supported onWrite(), which triggers when data is created, updated, or deleted in the Realtime Database. onCreate(), which triggers when new data is created in the Realtime Database. onUpdate(), which triggers when data is updated in the Realtime Database. onDelete(), which triggers when data is deleted from the Realtime Database. For any further detail: https://firebase.google.com/docs/database/extend-with-functions

- Cloud Run (fully managed)

Q13)

You are planning to migrate an app in GCP with these features: C# language Time activated It may need hours to complete elaboration A managed, effective and simple solution is required and preferred.

Which of the following solutions will you suggest (choose 1)?

- Compute Engine and Managed Instances
- ✅ Cloud Scheduler with App Engine .NET Flexible Environment Concepts with manual scaling

Explanation:-This is the only solution that supports the .NET environment and long term processing. Your instances with manual and basic scaling should run indefinitely, but there is no uptime guarantee. Hardware or software failures that cause early termination or frequent restarts can occur without warning and can take considerable time to resolve. All flexible instances are restarted on a weekly basis. During restarts, critical, backwards-compatible updates are automatically rolled out to the underlying operating system. Your application's image will remain the same across restarts. Cloud Scheduler can trigger the procedure on schedule A Compute Engine is not a managed solution C App Engine Standard Environment supports these programming languages: Python 2.7, Python 3.7 Java 8, Java 11 Node.js 8, Node.js 10 PHP 5.5, PHP 7.2, and PHP 7.3 Ruby 2.5 (beta) Go 1.9, Go 1.11, and Go 1.12 D Cloud Functions don't support C# E Cloud RUN don't support long term processing For any further detail: https://cloud.google.com/appengine/docs/flexible/dotnet/quickstart https://cloud.google.com/scheduler/

- Cloud Scheduler with App Engine Standard Environment
- Cloud Functions
- Cloud Run (fully managed)

Q14)

You are planning to migrate a set of microservices apps in GCP that are already organized and deployed in Containers.

You have been asked to find the best managed platform suited for these applications.

Scalability is a requirement but there are not sudden and high bursts of requests.

The applications are developed with different programming languages and tools and may be sessionless or may have to manage in memory sessions.

Which of the following solutions will you suggest (choose 1)?

- Compute Engine and Managed Instances
- ✅ App Engine Flexible Environment

Explanation:-This is the only platform that covers all the requirements. App Engine Flexible Environment gives a broad range of solutions with the only bound of using Containers, as in our case. Automatic scaling creates dynamic instances based on request rate, response latencies, and other application metrics. New instances start require more time compared to the Standard Environment, but it is clearly not an issue. Cloud Run is not suitable because it supports only sessionless applications. The other solutions are ot managed or don't support the runtimes required. For any further detail: https://cloud.google.com/appengine/docs/flexible/java/how-instances-are-managed

- App Engine Standard Environment
- Cloud Functions
- Cloud Run (fully managed)

Q15)

You need to migrate a .NET C # app to the Cloud with SQL Server.

You are asked which of the following solutions is the most portable and can guarantee a best platform independent solution in relation to the Cloud Environment.

Which of the following solutions will you suggest (choose 2)?

- Windows Compute Engine Instances
- ✅ App Engine Flexible Environment and DB Managed Service

Explanation:-The .NET runtime is the software stack responsible for installing your application code and its dependencies and running your application. It is already a way to make applications portable in different environments. In addition having the app containerized with Docker and orchestrated with Kubernetes organize even better cross-platform deployments. The other solutions simply don't have these features. For any further detail: https://cloud.google.com/appengine/docs/flexible/dotnet/runtime

- App Engine Standard Environment and DB Managed Service
- ✅ GKE and DB Managed Service

Explanation:-The .NET runtime is the software stack responsible for installing your application code and its dependencies and running your application. It is already a way to make applications portable in different environments. In addition having the app containerized with Docker and orchestrated with Kubernetes organize even better cross-platform deployments. The other solutions simply don't have these features. For any further detail: https://cloud.google.com/appengine/docs/flexible/dotnet/runtime

---

Q16)

You have 2 different Customers that want to transfer - deploy their application/code to the Cloud. Both of them want to use managed services.

The first one is a big Enterprise Company that has legacy and modern Applications.

It uses a famous ERP and SQL Databases.

The The second one is a startup Company that has microservices Functions and noSQL Databases.

Which of the following solutions will you suggest for the Enterprise Company (choose 1)?

- Compute Engine and Managed Instances
- ✅ App Engine Flexible Environment and Managed DB Services

Explanation:-App Engine Flexible Environment gives all the technological capabilities with the only bound of using Containers. Your Customer may deploy in the Cloud any application with minimum effort and great advantages: scalability, economy, flexibility. For any further detail: https://cloud.google.com/appengine/docs/flexible/java/how-instances-are-managed

- App Engine Standard Environment and Managed DB Services
- Cloud Functions and Managed DB Services
- Cloud Run (fully managed)

Q17)

You have 2 different Customers that want to transfer - deploy their application/code to the Cloud. Both of them want to use managed services.

The first one is a big Enterprise Company that has legacy and modern Applications.

It uses a famous ERP and SQL Databases.

The The second one is a startup Company that has microservices Functions and a noSQL Databases.

Tha Management doesn't want fixed costs.

Which of the following solutions will you suggest for the Startup Company (choose 2)?

- Compute Engine and Managed Instances
- App Engine Flexible Environment and Managed DB Services
- App Engine Standard Environment and Managed DB Services
- ✅ Cloud Functions and Cloud Datastore (Firebase Edition)

Explanation:-These are the solutions that minimize costs, but you have to design backend Apps in a modern way and, above all, stateless. Cloud Datastore is a fully NoSQL document database with automatic scaling, high performance, and ease of application development. Firestore, the newest version of Datastore, makes all queries strongly consistent. Cloud Functions is a serverless execution environment for building simple microservices apps and connecting cloud services. Cloud Run (fully managed) is a serverless, fully managed compute platform that automatically scales stateless containers. For any further detail: https://cloud.google.com/datastore/docs/concepts/overview https://cloud.google.com/functions/docs/concepts/overview https://cloud.google.com/run/

- ✅ Cloud Run (fully managed) and Cloud Datastore (Firebase Edition)

Explanation:-These are the solutions that minimize costs, but you have to design backend Apps in a modern way and, above all, stateless. Cloud Datastore is a fully NoSQL document database with automatic scaling, high performance, and ease of application development. Firestore, the newest version of Datastore, makes all queries strongly consistent. Cloud Functions is a serverless execution environment for building simple microservices apps and connecting cloud services. Cloud Run (fully managed) is a serverless, fully managed compute platform that automatically scales stateless containers. For any further detail: https://cloud.google.com/datastore/docs/concepts/overview https://cloud.google.com/functions/docs/concepts/overview https://cloud.google.com/run/

Q18)

Your company is expanding its business worldwide. You have been asked which of the GCP services that are currently used are already accessible all over the world and which ones are linked to their region.

Below is a list of services.

Which are the Global ones (pick 2)?

- L3 Load Balancer
- ✅ L7 Load Balancer

Explanation:-L7 Load Balancer is the HTTP(S) load balancer. It is distributed globally, so as to minimize latency between clients and the load balancer. A: Network Load Balancing (L3) terminates TLS on backends that are located in specific regions .

- Data Prep
- ✅ Cloud Spanner

Explanation:-Cloud Spanner is the first scalable, enterprise-grade, globally-distributed, and strongly consistent database service built for the cloud specifically to combine the benefits of relational database structure with non-relational horizontal scale. C: Cloud Dataproc is a managed Spark and Hadoop service that lets you take advantage of open source data tools for batch processing, querying, streaming, and machine learning. Cloud Dataproc automation helps you create REGIONAL clusters quickly, manage them easily, and save money by turning clusters off when don't needed. It may support "global" endpoint is a special multi-region namespace but it is Regional. E: Kubernetes Engine environment consists of multiple machines (specifically, Google Compute Engine instances) grouped together to form a cluster that usually runs in a single compute zone. Regional clusters increase the availability of both a cluster's control plane (master) and its nodes by replicating them across multiple zones of a region. This provides the advantages of multi-zonal clusters. For any further detail: https://cloud.google.com/load-balancing/docs/choosing-load-balancer https://cloud.google.com/load-balancing/docs/network/ https://cloud.google.com/dataproc/docs/concepts/regional-endpoints https://cloud.google.com/dataproc/docs/concepts/overview https://cloud.google.com/kubernetes-engine/docs/concepts/kubernetes-engine-overview

- GKE

Q19)

Your direct Manager asked you for advice: he needs shared disk storage for his team.

It has to deal with any types of objects and executables, be simple and easy to use and manage, not expensive and secure.

He cannot make up his mind among the following products.

Which one is the most suitable for his requirements?

- Local SSD
- Regional persistent disk and regional SSD persistent disk
- ✅ Filestore

Explanation:-Filestore You can easily mount Cloud Filestore file shares on Compute Engine instances, enabling visual effects artists to collaborate on the same file share. As rendering workflows typically run across fleets (“render farms”) of compute machines, all of which mount a shared filesystem, Filestore and Compute Engine can scale to meet your job’s rendering needs. Google Cloud CDN leverages Google's globally distributed edge points of presence to accelerate content delivery for websites and applications served out of Compute Engine and Cloud Storage. For any further detail: https://cloud.google.com/compute/docs/disks/ https://cloud.google.com/cdn/docs/overview https://cloud.google.com/storage/docs/ https://cloud.google.com/filestore/

- Cloud Storage
- Cloud Storage and Cloud CDN

Q20)

Your company has made a business acquisition and needs to merge the portal's IT services. It is necessary to integrate procedures that are written in different languages ??and then start a development project for the creation of a flexible and managed system that offers, in a simple way, the possibility of creating, maintaining, deploying and offering in a scalable way services from code developed with various different technologies.

Which of the following is the ideal combination of GCP products to meet the required requirements?

- Cloud Functions and https
- ✅ App Engine and Cloud Endpoints

Explanation:-App Engine, both Standard then Flex Edition, are specially suited for Building Microservices. In an App Engine Project you can use any mix of standard and flexible environment services, written in any language. In addition with Cloud Endpoints it is possible to deploy, protect, and monitor your APIs. Using an OpenAPI Specification or API frameworks, Cloud Endpoints gives tools for API development and provides insight with Stackdriver Monitoring, Trace, and Logging. Cloud Functions with https endpoint is not enough for enterprise integrated projects. C: Load Balancers e Cloud Armor is completely wrong. Cloud Armor protects against malicious attacks, so it is out of our scope. D: Google Cloud CDN leverages Google's globally distributed edge points of presence to accelerate content delivery for websites and applications served out of Compute Engine and Cloud Storage. The lifecycle management of Cloud Storage makes it easy to manage deletion after a given period of time. configure a Cloud Storage bucket to host a static website for a domain you own. Static web pages can contain client-side technologies such as HTML, CSS, and JavaScript. They cannot contain dynamic content F: Firebase Hosting is suitable for mobile and web applications; it is a production-grade web content hosting for developers for serving both static and dynamic content to a global CDN. Out of our scope. For any further detail: https://cloud.google.com/appengine/docs/ https://cloud.google.com/endpoints/ https://firebase.google.com/docs/hosting/

- Load Balancers and Cloud Armor
- Cloud Storage and Cloud CDN
- Firebase Hosting

Q21)

Your customer needs to collect and analyze a very large number of data from industrial machinery.

He chose to use Bigtable because he needs to access data with latis of milliseconds on very large numbers, with categories of information that can greatly vary.

He asked for your advice to organize the structure and especially the access key, since he was informed that this is a crucial decision.

Which kind of a Row Key do you choose?

- Timestamp plus DeviceID
- Sequence Number
- ✅ Combination of DeviceID plus timestamp

Explanation:-With Bigtable it is useful to include multiple identifiers in row key. When row key includes multiple values, it's especially important to have a clear understanding of data usage and to create a row key that still makes it possible to retrieve a well-defined range of rows. Otherwise the queries will require a table scan, which is much slower than retrieving specific rows. A is wrong because most writes would be pushed onto a single node. For the same reason, avoid placing a timestamp at the start of the row key B is wrong because in this way you push most of your traffic to a small number of nodes. A safer approach is to use a reversed version of the user's numeric ID, which spreads traffic more evenly across all of the nodes for your Cloud Bigtable table. D is wrong because it is not recommended to use a hash. Hashed row keys make it very difficult to troubleshoot issues with Cloud Bigtable, because hashed row keys are effectively meaningless. For example, if your row key is a hash of the device ID, it will be difficult or impossible to find out what user ID is tied to the row key. Bigtable DBs are stored in sparse, distributed and multidimensional tables. Each record is identified by row key, column key, and timestamp, and are in order by row key. The row range for a table is dynamically partitioned. Each row range is called a tablet, which is the unit of distribution and load balancing. Column keys are grouped into sets called column families. All data stored in the same column family are usually of the same type. For any further detail: https://cloud.google.com/bigtable/docs/schema-design https://www.sciencedirect.com/topics/computer-science/google-bigtable

- Hash of a string ID

Q22)

Your company is planning to migrate to GCP. Migration of SQL databases needs to be addressed. Your company is international and therefore needs a global database for integrated ERP and local databases for specific applications for each country.

All data must be integrable and will be stored into an Analytics System, too. All services must be fully managed.

The on-premises Systems used to have primary keys represented by a sequence, a progressive number.

Which GCP products do you choose (pick 3)?

- ✅ Cloud SQL

Explanation:-You will advise to use the main 3 SQL GCP Database. Cloud SQL for local Databases: not expensive, reliable, regional, secure. Spanner as a global SQL Database for storing Data of international Applications. BigQuery for Analytics. C and E are wrong because they are noSQL Database and so, they address completely different requirements. For any further detail: https://cloud.google.com/bigquery/docs/ https://cloud.google.com/sql/docs/sqlserver/ https://cloud.google.com/spanner/ https://cloud.google.com/sql/docs/

- ✅ BigQuery

Explanation:-You will advise to use the main 3 SQL GCP Database. Cloud SQL for local Databases: not expensive, reliable, regional, secure. Spanner as a global SQL Database for storing Data of international Applications. BigQuery for Analytics. C and E are wrong because they are noSQL Database and so, they address completely different requirements. For any further detail: https://cloud.google.com/bigquery/docs/ https://cloud.google.com/sql/docs/sqlserver/ https://cloud.google.com/spanner/ https://cloud.google.com/sql/docs/

- Firestore

Explanation:-You will advise to use the main 3 SQL GCP Database. Cloud SQL for local Databases: not expensive, reliable, regional, secure. Spanner as a global SQL Database for storing Data of international Applications. BigQuery for Analytics. C and E are wrong because they are noSQL Database and so, they address completely different requirements. For any further detail: https://cloud.google.com/bigquery/docs/ https://cloud.google.com/sql/docs/sqlserver/ https://cloud.google.com/spanner/ https://cloud.google.com/sql/docs/

- ✅ Spanner
- Bigtable

Q23)

You are designing an IoT system that collects and processes a large amount of data from different devices for a Smart City project (cameras, sensors, control devices, etc.).

Are you wondering which is the best GCP product and the data scheme to use, considering that:

The information of each type of device may be different You are interested in immediate identification by area, time, device type You need access times with latencies in milliseconds The project has a limited geographical scope

You need to be able to group and identify the data within each record in various ways

Which of the following products do you choose?

- Cloud Datastore
- Cloud Spanner
- ✅ Cloud Bigtable

Explanation:-Bigtable is the perfect choice. Bigtable is suitable for ad tech, fintech, and IoT and offers consistent sub-10ms latency. Replication provides higher availability, higher durability, and resilience in the face of zonal failures. Cloud Bigtable is designed with a storage engine for machine learning applications and provides easy integration with open source big data tools Regarding the other answer, for a flexible schema you need a noSQL Database. So, Datastore or Bigtable. All the others are SQL DBs. Datastore cannot supply latencies in milliseconds with petabytes of data. Bigtable can, definitely. For any further detail: https://cloud.google.com/bigtable/docs/schema-design https://cloud.google.com/bigtable/

- BigQuery
- Cloud SQL

Q24)

For a project of yours, it is required a database with a non-rigid, high-performance schema that can easily manage Customers, Orders and Invoices relationships;

in other words, you need to deal with hierarchically structured objects and you are looking for an economically convenient solution. In addition transactions, with serializable isolation enforcement are required.

Which of the following products do you choose?

- ✅ Cloud Datastore

Explanation:-Datastore manages relationships between entities (records), in a hierarchically structured space similar to the directory structure of a file system. When you create an entity, you can optionally designate another entity as its parent; the new entity is a child of the parent entity. An entity without a parent is a root entity. A transaction is a set of Datastore operations on one or more entities in up to 25 entity groups. Each transaction is guaranteed to be atomic, which means that transactions are never partially applied. Either all of the operations in the transaction are applied, or none of them are applied. Regarding the other answer, for a flexible schema you need a noSQL Database. So, Datastore or Bigtable. All

the others are SQL DBs. Datastore is the one that can manage transactions, even with serializable isolation enforcement. Bigtable don't manage transactions. For any further detail: https://cloud.google.com/datastore/docs/concepts/entities#ancestor_paths https://cloud.google.com/datastore/docs/concepts/cloud-datastore-transactions

- Cloud Spanner
- Cloud Bigtable
- BigQuery
- Cloud SQL

Q25)

You are responsible for planning the migration to GCP of an important application that works with Oracle Database. A horizontally scalable and globally functioning SQL database is required.

Which service is better to use and which type of schema migration is recommended?

- Cloud SQL with no schema migration
- Cloud SQL with sequential primary keys migration
- Cloud Spanner with no schema migration
- ✅ Cloud Spanner with sequential primary keys migration

Explanation:-The requirements point to an SQL Database that is global and distributed with synchronized replicas and shards in multiple Servers: Cloud Spanner. The risk of hotspotting with synchronized replicas needs to be addressed; that is updates that are not distributed among multiple servers. So it is necessary to be careful not to create hotspots with the choice of your primary key. For example, if you insert records with a monotonically increasing integer as the key, you'll always insert at the end of your key space. This is undesirable because Cloud Spanner divides data among servers by key ranges, which means your inserts will be directed at a single server, creating a hotspot. The techniques that can spread the load across multiple servers and avoid hotspots: Hash the key and store it in a column. Use the hash column (or the hash column and the unique key columns together) as the primary key. Swap the order of the columns in the primary key. Use a Universally Unique Identifier (UUID). Version 4 UUID is recommended, because it uses random values in the high-order bits. Don't use a UUID algorithm (such as version 1 UUID) that stores the timestamp in the high order bits. Bit-reverse sequential values. For any further detail: https://cloud.google.com/spanner/docs/schema-and-data-model

---

Q26)

You are looking for a SQL system to integrate and query both historical and production data.

The data must be organized in complex structures.

In particular, it is necessary to store orders and invoices in a denormalized and complete manner with the header and detail within the same structure.

Which of the following products do you choose?

- Cloud Datastore
- Cloud Spanner
- Cloud Bigtable
- ✅ BigQuery

Explanation:-BigQuery is an OLAP engine. So, it is far better, even if it can manage normalised data and joins, to have denormalized information. In addition BigQuery can manage nested and repeated columns and structures, as required. BigQuery is not a Database but an enterprise, serverless, highly scalable, and cost-effective cloud data data warehouse that solves this problem by enabling super-fast SQL queries using the processing power of Google's infrastructure. It can quickly analyze gigabytes to petabytes of data using ANSI SQL. For any further detail: https://cloud.google.com/bigquery/what-is-bigquery https://cloud.google.com/bigquery/docs/nested-repeated

- Cloud SQL

Q27)

Your Company is planning to use BigQuery as a cloud data warehouse.

You asked for a list of the main features and advice related to this service, and, as soon as you got it, you realized that some items were wrong.

Which ones?

- Partitioned and clustered tables lets you spend less when you query the specific data and use the LIMIT statement
- You can set a table's expiration time
- ✅ The LIMIT statement saves you money

Explanation:-If you use partitioned and clustered tables and the where clause limits data in the partition, the LIMIT clause is a money saver B is true; you can set a TTL (time to live) C is wrong; with LIMIT you still pay for the table scan. D is wrong. BigQuery can manage normalised data and joins but performs better with denormalized information. E is true; Bigquery bills are related to the amount of processed data F is wrong; Bigquery uses Colossus directly. In addition, BigQuery uses Bigtable for its streaming engine, and Spanner for its metadata and query result preview. For any further detail: https://cloud.google.com/bigquery/what-is-bigquery https://cloud.google.com/bigquery/docs/managing-tables https://polleyg.dev/posts/data-engineering-tips/

- ✅ BigQuery wants denormalized data because it cannot handle joins

Explanation:-If you use partitioned and clustered tables and the where clause limits data in the partition, the LIMIT clause is a money saver B is true; you can set a TTL (time to live) C is wrong; with LIMIT you still pay for the table scan. D is wrong. BigQuery can manage normalised data and joins but performs better with denormalized information. E is true; Bigquery bills are related to the amount of processed data F is wrong; Bigquery uses Colossus directly. In addition, BigQuery uses Bigtable for its streaming engine, and Spanner for its metadata and query result preview. For any further detail: https://cloud.google.com/bigquery/what-is-bigquery https://cloud.google.com/bigquery/docs/managing-tables https://polleyg.dev/posts/data-engineering-tips/

- Avoid the use of SELECT * : more data more expenses
- ✅ BigQuery's storage layer is GCS

Explanation:-If you use partitioned and clustered tables and the where clause limits data in the partition, the LIMIT clause is a money saver B is true; you can set a TTL (time to live) C is wrong; with LIMIT you still pay for the table scan. D is wrong. BigQuery can manage normalised data and joins but performs better with denormalized information. E is true; Bigquery bills are related to the amount of processed data F is wrong; Bigquery

uses Colossus directly. In addition, BigQuery uses Bigtable for its streaming engine, and Spanner for its metadata and query result preview. For any further detail: https://cloud.google.com/bigquery/what-is-bigquery https://cloud.google.com/bigquery/docs/managing-tables https://polleyg.dev/posts/data-engineering-tips/

Q28)

You are designing an app that allows users to upload images and videos to the Cloud from the web and mobile interfaces.

The users are not allowed to permanent permissions for uploading objects, but the application must provide them with the ability to carry on these tasks only when required.

Which of the following is the best technique to use?

- Create a "dummy" user with the necessary rights that is used for the purpose by the application
- ✅ Use of signed URLs

Explanation:-A signed URL is a simple, clean and economic solution: they give time-limited resource access to anyone in possession of the URL, regardless of whether they have a Google account. A signed URL is a URL that provides limited permission and time to make a request. Signed URLs contain authentication information in their query string, allowing users without credentials to perform specific actions on a resource. When you generate a signed URL, you specify a user or service account which must have sufficient permission to make the request that the signed URL will make. After you generate a signed URL, anyone who possesses it can use the signed URL to perform specified actions, such as reading an object, within a specified period of time. For any further detail: https://cloud.google.com/storage/docs/access-control/signed-urls

- Use the application credentials
- Provide the user with a scheduled temporary permit

Q29)

The management of your company asked you to design a static website aimed at hosting your company's product sheets.

The website must be cheap and simple to set up and manage.

The site must have the corporate domain, but it does not need to be served through HTTPS.

What is the optimal and more economic solution among the following?

- ✅ Cloud Storage

Explanation:-A Cloud Storage bucket can host a static website for a domain you own. Static web pages can contain client-side technologies such as HTML, CSS, and JavaScript. They cannot contain dynamic content such as server-side scripts like PHP. There is no additional cost beyond storage. It is possible only to serve using direct URIs such as https://storage.googleapis.com/my-bucket/my-object because when hosting a static website using a CNAME redirect, Cloud Storage only supports HTTP. In the case HTTPS serving is required, you can: Set up a load balancer. Use a third-party Content Delivery Network with Cloud Storage. Serve your static website content from Firebase Hosting instead of Cloud Storage. For any further detail: https://cloud.google.com/storage/docs/hosting-static-website

- Compute Engine
- App Engine
- Cloud Functions
- Firebase Hosting

Q30)

You have two customers with a very similar problem: the transformation of applications from stateful to stateless, e regarding session information (user data, activities carried out, current state of the work session).

The first customer is a large enterprise with a wide spread and traffic for its services that wants to maintain high performance.

The Second Client is a startup company, less traffic and a lot of attention to savings.

All these applications use SQL databases.

Which solution do you choose for the enterprise customer?

- Cloud SQL
- BigQuery
- Cloud Datastore
- ✅ Cloud Memorystore

Explanation:-The best solution is Cloud Memorystore, that is a noSQL Database in memory. Cloud Memorystore for Redis is a fully managed in-memory data store service built on scalable, secure, and highly available infrastructure managed by Google. Cloud Memorystore caches data and provides sub-millisecond response times. Cloud Datastore may be suitable but with less performances. It is less expensive, too. The other soluzions are SQL ones, and so they are not at all suitable because the scheme is not flexible. For any further detail: https://cloud.google.com/memorystore/

Q31)

You have two customers with a very similar problem: the transformation of applications from stateful to stateless, e regarding session information (user data, activities carried out, current state of the work session).

The first customer is a large enterprise with a wide spread and traffic for its services that wants to maintain high performance.

The Second Client is a startup company, less traffic and a lot of attention to savings.

All these applications use SQL databases.

Which solution do you choose for the startup customer?

- BigQuery
- Cloud Datastore
- ✅ Cloud Memorystore

Explanation:-The best solution is Cloud Memorystore, that is a noSQL Database in memory. But is too expensive for out startup. Cloud Datastore is

suitable but with less performances. It is less expensive, too. The other soluzions are SQL ones, and so they are not at all suitable because the scheme is not flexible. For any further detail: https://cloud.google.com/datastore

- Cloud SQL

Q32)

You have an ERP System that works with SQL Database. This information are used for commercial transactions and until the monthly accounting closure the data must be available for reviews and their consistency must always be guaranteed, so transactions must be guaranteed according to the ACID rules. Once closed, the data cannot be changed anymore.

What organization can you recommend to your client:

- Keep all the data in the ERP also for querying purposes
- ✅ After the monthly closings, transfer the data to BigQuery with SQL, from which to query the historical and production data

Explanation:-BigQuery is an analytics system and it can directly query and import data from SQL Databases with federated query. With this method it is possible to transfer data and query a great amount of historical data. BigQuery is not an ACID Database, so, when the data is updatable, it has to be stored in an SQL Database. Acid is an acronym that means: Atomic All operations in a transaction succeed or every operation is rolled back. Consistent On the completion of a transaction, the database is structurally sound. Isolated Transactions do not contend with one another. Contentious access to data is moderated by the database so that transactions appear to run sequentially. Durable The results of applying a transaction are permanent, even in the presence of failures. For any further detail: https://cloud.google.com/bigquery/docs/cloud-sql-federated-queries

- After the monthly closings, transfer the data to BigTable with SQL, from which to query the historical and production data
- After the monthly closings, transfer the data with SQL to Datastore, from which to query the historical and production data

---

Q33)

You have been asked to identify the most suitable storage solution for registering files that: are transmitted from external locations. are read and processed overnight by a procedure that runs in a single VM are loaded into the centralized database.

A fast and secure processing is required.

What solutions do you plan to adopt to store the files and programs described?

- Local SSD
- Zonal SSD persistent disks
- ✅ Regional SSD persistent disks

Explanation:-SSD Disks are fast and Regional persistent disks provide durable storage and replication of data between two zones in the same region. So data and processing can be secured. A is wrong because Local SSD are directly attached to the VM and so they are transient and data can be lost. B is wrong because data is not replicated D is wrong because RAM Disks are in memory so data can be lost. E is wrong because Filestore: is a high performance file storage for Google Cloud users and not for a single VM. For any further detail: https://cloud.google.com/compute/docs/disks/ https://cloud.google.com/compute/docs/disks/local-ssd

- RAM Disks
- Filestore

Q34)

With Cloud Storage you may have different classes and it is possible to pass from one class to another. But some transitions are not allowed.

Which one of the following is not possible?

- Regional to nearline
- Multiregional to coldline
- ✅ Regional to multiregional

Explanation:-When you create a bucket you have to declare if it will be either regional or multiregional, You cannot change afterwards. All the other transitions are allowed. For any further detail: https://cloud.google.com/storage-transfer/docs/overview

- Nearline to coldline

Q35)

A customer of yours is planning to migrate their infrastructure to the Cloud.

It is an international Company that has the requirement to transfer web and mobule applications in a short time, and only afterwards gradually optimize and re-engineer them.

Se, they want to perform a lift and shift migration and therefore improve the systems according to the strangler pattern.

Which of the following strategies is the most advisable for your Customer?

- Use Cloud Tasks
- Create 2 different project with App Engine standard edition
- ✅ Use Endpoints for OpenAPI Use Cloud Functions

Explanation:-You can use Endpoints for OpenAPI as an interface so you can gradually replace specific pieces of functionality served with legacy apps in Compute Engine with the new software developed, for example, with serverless technologies, the ones that you do prefer. Endpoints is an API management system that helps to secure, monitor, analyze, and set quotas to backends. Moreover, after you deploy your API to Endpoints, you can use the Cloud Endpoints Portal to create a developer portal, a website that users of your API can access to view documentation and interact with your API. A is wrong because Cloud Tasks is an asynchronous task execution service B is wrong because it is not related only to App Engine D and E are wrong because you can use all the technologies you need and prefer, not only Cloud Functions and GKE For any further detail: https://cloud.google.com/endpoints/docs/

- Use GKE

Q36)

You are the leader of a development group that is migrating some applications to the Cloud and who has asked you how to set up a local work environment.

In the company they need to use specific development tools that are installed on the Clients.

Which of these tips would you provide?

- Develop locally by making remote calls to services with credentials within the code
- Develop remotely with the setup of a VM with Compute Engine and a Development Disk Image
- Use Cloud Shell
- ✅ Install Cloud SDK and use Service Accounts

Explanation:-Cloud SDK is scoped to this very aim and Service Accounts are the Google recommended practice for authorization For any further detail: https://cloud.google.com/sdk/docs/quickstarts https://cloud.google.com/shell/docs/using-cloud-shell

---

Q37)

You have been asked for advice from a person who has never operated on GCP and who is a little puzzled.

He asked you a series of doubts, based on the indications they provided.

Which of these statements are correct (pick 3) ?

- Projects, Billing accounts and Service accounts are resources and identities
- ✅ Before creating a project is advisable to setup a Resource Hierarchy

Explanation:-Resource Hierarchy is an organization very useful when user and resources in GCP are increasing and need to be managed in a clear and economic way. C is wrong because Organization are on top of the Hierarchy D is correct as you can easily see in the following picture E is correct because it is very handy to set once policies shared in the related section of the Hierarchy For any further detail: https://cloud.google.com/iam/docs/resource-hierarchy-access-control

- You can have a project related to many organizations
- ✅ An organization is the root of the resource hierarchy

Explanation:-Resource Hierarchy is an organization very useful when user and resources in GCP are increasing and need to be managed in a clear and economic way. C is wrong because Organization are on top of the Hierarchy D is correct as you can easily see in the following picture E is correct because it is very handy to set once policies shared in the related section of the Hierarchy For any further detail: https://cloud.google.com/iam/docs/resource-hierarchy-access-control

- ✅ Multiple policies can exists for a folder or project and they propagate down the structure

Explanation:-Resource Hierarchy is an organization very useful when user and resources in GCP are increasing and need to be managed in a clear and economic way. C is wrong because Organization are on top of the Hierarchy D is correct as you can easily see in the following picture E is correct because it is very handy to set once policies shared in the related section of the Hierarchy For any further detail: https://cloud.google.com/iam/docs/resource-hierarchy-access-control

---

Q38)

You are the leader of a development group; you want to start using Continuous Integration and Deployment Techniques and you care to organize procedure in the best way.

In your company you are not allowed to publish code in public or not internally certified Sites.

Where will the code developed by your team be stored and shared?

- Cloud Storage with versioned Objects
- Github
- ✅ Cloud Source Repository

Explanation:-Google Cloud Source Repositories are private, fully featured, scalable Git repositories hosted on Google Cloud Platform. Git is a program that monitors files and tracks changes. A Git repository: Tracks any updates Registers a history May trigger actions A is wrong because Cloud Storage only registers the complete versions of the files. B is wrong because Github is a Git repository but not private or not internally certified. D is wrong because AppEngine is not a Git repository and Blue green is a kind of deployment, not source integration tool. For any further detail: https://cloud.google.com/docs/ci-cd/

- AppEngine and Blue green Integration

Q39)

You are the leader of a development group; you want to start using Continuous Integration and Deployment Techniques and you care to organize procedure in the best way.

In your company the new trend is to deploy apps and services within containers.

The idea is to use Kubernetes.

You want to start the deployment as soon as new Source is committed.

Which product is the best suitable one for creating Docker images from code?

- ✅ Cloud Build

Explanation:-Cloud Build can define workflows for building, testing, and deploying across multiple environments such as VMs, serverless, Kubernetes, or Firebase. B is wrong because Cloud Code is an integrated set of tools to help write, deploy, and debug cloud-native applications. It as extensions to IDEs such as Visual Studio Code and IntelliJ are provided to let rapidly iterate, debug, and deploy code to Kubernetes. C is wrong because Cloud Tasks is an asynchronous task execution service that encode and execute Tasks using Queues. D is wrong because Cloud Repositories are Git source repos. E is wrong because Cloud Run is a serverless platform for containerized applications. For any further detail: https://cloud.google.com/docs/ci-cd/ https://cloud.google.com/cloud-build/docs/quickstart-docker https://cloud.google.com/source-repositories/docs/quickstart-triggering-builds-with-source-repositories

- Cloud Code
- Cloud Tasks

- Cloud Repositories
- Cloud Run

Q40)

Your team is developing an app with a serverless backend.

A developer had an idea: automate the Unit Testing process in order to increase the efficiency and quality of the software produced.

You have already imported Cloud Build triggers for automatic deployment.

What have you to do for automatic unit tests (pick 3)?

- Nothing. Cloud Build has everything
- ✅ Use a mocking framework for simulation of external actions

Explanation:-A CI/CD platform such as Cloud Build can be configured to run your tests on an ongoing basis. Continuous testing helps ensure that your code continues to work as intended and that your dependencies remain up-to-date. You will need a Testing frameworks, such as Mocha, that provides an execution environment for automated tests and for getting results back (Log and Lo sinks help). In addition Mocking frameworks let you mock external dependencies. An external dependency is a dependency that your function relies on that isn't a part of your function's code. Finally you can configure CI/CD pipelines with Cloud Build to automatically test. For any further detail: https://cloud.google.com/functions/docs/testing/test-cicd

- ✅ Use a test framework for preparing unit tests and return results

Explanation:-A CI/CD platform such as Cloud Build can be configured to run your tests on an ongoing basis. Continuous testing helps ensure that your code continues to work as intended and that your dependencies remain up-to-date. You will need a Testing frameworks, such as Mocha, that provides an execution environment for automated tests and for getting results back (Log and Lo sinks help). In addition Mocking frameworks let you mock external dependencies. An external dependency is a dependency that your function relies on that isn't a part of your function's code. Finally you can configure CI/CD pipelines with Cloud Build to automatically test. For any further detail: https://cloud.google.com/functions/docs/testing/test-cicd

- Use Cloud RUN to prepare the test and automate it
- ✅ Use Cloud Build Triggers for automating the pipeline

Explanation:-A CI/CD platform such as Cloud Build can be configured to run your tests on an ongoing basis. Continuous testing helps ensure that your code continues to work as intended and that your dependencies remain up-to-date. You will need a Testing frameworks, such as Mocha, that provides an execution environment for automated tests and for getting results back (Log and Lo sinks help). In addition Mocking frameworks let you mock external dependencies. An external dependency is a dependency that your function relies on that isn't a part of your function's code. Finally you can configure CI/CD pipelines with Cloud Build to automatically test. For any further detail: https://cloud.google.com/functions/docs/testing/test-cicd

---

Q41)

Your Team is ready to start the CI/CD adventure.

You want to extend your Pipelines to servers on-premises or other cloud providers.

Which additional product is advisable to use with Cloud Build, Container Registry?

- Data Prep
- ✅ Spinnaker

Explanation:-Spinnaker is an open source, multi-cloud continuous delivery platform for releasing software changes. Its main functions:
*) Automate deployment pipelines that run integration and system tests, spin up and down server groups, and monitor your rollouts via git events, Jenkins, Travis CI, Docker, CRON, or other Spinnaker pipelines.
*)Deploy across multiple cloud providers including AWS EC2, Kubernetes, Google Compute Engine, Google Kubernetes Engine, Google App Engine, Microsoft Azure, Openstack, Cloud Foundry, and Oracle Cloud Infrastructure
*)Create and deploy immutable images For any further detail: https://cloud.google.com/solutions/continuous-delivery/ https://www.spinnaker.io/

- ISTIO
- Cloud RUN
- Cloud Source

Q42)

Your team is developing an app with a serverless backend. With your team you automated the Unit Testing process in order to increase the efficiency and quality of the software produced.

You choose Mocha as a Testing frameworks that provides an execution environment for automated tests and for getting results back into Stackdriver Logs.

Now you have to better organize this last step.

Which is the best choice (pick 1)?

- Create a function to browse log and extract test results
- ✅ Create a log sink to Storage for the test results

Explanation:-You may run each step in the Cloud Build pipeline in a separate container, so the logs for each step are isolated and easily retrievable. You set In addition, every time a log entry is registered, if query matches the log entry writes a copy of the log entry to the sink's export destination. Since exporting happens for new log entries only, you cannot export log entries that Logging received before your sink was created. For any further detail: https://cloud.google.com/functions/docs/testing/test-cicd https://cloud.google.com/logging/docs/export/ https://cloud.google.com/functions/docs/testing/test-basics

- Export logs to files and use Data Prep
- Use a ML algorithm
- Associate a Pub/Sub to the logs

Q43)

Your team is developing an app for IoT data integration, You are ready to start production but you are thinking about all the things that could possibly fail.

One of those is the possibility of a high number of transmissions that could make the system unresponsive.

So, you decide to perform a special kind of test.

What is the best among the following ones?

- System tests
- Performance Test
- Unit Test
- ✅ Load Test

Explanation:-Load testing is a type of software testing to understand the behavior of the application under a specific expected load, under both normal and at peak conditions. It is a part of performance testing. The most used tool for this kind of test is called Locust, because it simulates a Locusts' attack. For any further detail: https://cloud.google.com/solutions/distributed-load-testing-using-gke

- Integration test

Q44)

Your team is developing an app that will be deployed with App Engine. You completed Unit Tests.

This App uses a lot of GCP and other providers Services.

You are ready to start production but you are thinking about all the things that could possibly fail.

You want to perform an Integration test.

Which of the following actions can your team execute?

- Test locally with shims
- Use Split health checks
- Create a new version with the --no-promote flag and perform tests
- Use the truncated exponential backoff for periodic polling if asynchronous operations
- ✅ All of the above

Explanation:-Integration tests validate interaction between sections of your code, and can be used to test a usage of other Google Cloud services such as Databases or external services. The primary difference between unit tests and integration tests for Cloud Functions is that integration tests involve less mocking than unit tests. Integration tests should trigger and respond to actual Cloud events such as HTTP requests, Pub/Sub messages, or Storage object changes. You can run integration tests locally using a shim. Some asynchronous operations may require periodic polling. The truncated exponential backoff polling strategy effectively minimizes both the underlying operation's latency as well as the polling frequency required. Truncated exponential backoff is a standard error handling strategy for network applications in which a client periodically retries a failed request with increasing delays between requests. In App Engine, by default, split health checks are enabled. You can use periodic health check requests to confirm that a VM instance has been successfully deployed, and to check that a running instance maintains a healthy status. Each health check must be answered within a specified time interval. Before configuring a new version to receive traffic, you can test it on App Engine. Traffic migration switches the request routing between the versions within a service of your application, moving traffic from one or more versions to a single new version. For any further detail: https://cloud.google.com/functions/docs/testing/test-basics
https://cloud.google.com/appengine/docs/flexible/python/testing-and-deploying-your-app

---

Q45)

Your team is developing an app that will be deployed with App Engine.

You completed Unit, Integration and Load Tests.

You gathered a nice documentation of all the activities executed.

What else is advisable to do?

- Nothing, GCP will run the app in the best way
- Set Stackdriver alerts on thresholds
- ✅ Use Stackdriver Performance Tools

Explanation:-A very good way to maintain optimal performance is to use Stackdriver Profiler and Stackdriver Trace. Stackdriver Profiler is a statistical, low-overhead profiler that continuously gathers CPU usage and memory-allocation information from your production applications. So it meets our requirements because it helps to identify the parts of the application consuming the most resources, and he performance characteristics of the code. Stackdriver Trace is a tracing system that collects latency data and displays it in near real time in the Google Cloud Platform Console. A is wrong because if you misconfigured some services performances CAN go bad. B is correct because Stackdriver alerts on thresholds are used for VM, so Compute Engine D is wrong because Cloud Armor is a distributed service to provide built-in defenses against infrastructure DDoS attacks. For any further detail: https://cloud.google.com/profiler/ https://cloud.google.com/trace/

- Use Cloud Armor

Q46)

The management asked you, as project leader of the development, to prepare a plan with the organizational proposals for the migration of corporate apps to the cloud, both as development projects and as an operational strategy.

What do you propose for the new organization of development?

- ✅ Create small groups with scrum master

Explanation:-In order to improve and modernize software development it is advisable to adopt Agile. Agile software development organized requirements and solutions through the collaborative effort of self-organizing and cross-functional teams and their customer/end user. It aims at adaptive planning, evolutionary development, early delivery, and continual improvement. Scrum is one of the most known agile process framework for managing complex knowledge work. For any further detail: https://www.scrumguides.org/scrum-guide.html

- No changes: keep on consolidated methods
- Decentralize development and use offshore development

Maintain the current organization and modernize the deployment

---

Q47)

A company asked you to plan a systems migration and a new technological architecture for the Cloud.

You have been asked to indicate a series of patterns that can allow you to obtain greater efficiency in processing large amounts of data.

Which of these options do you suggest (pick 3)?

Throttling

Cache-Aside

Explanation:-Cache-Aside Load data on demand into a cache from a data store. This can improve performance and also helps to maintain consistency between data held in the cache and data in the underlying data store. If an application updates information, it can follow the write-through strategy by making the modification to the data store, and by invalidating the corresponding item in the cache. Materialized View Prepopulated views over the data in one or more tables when the data isn't ideally structured for required query operations. This can improve querying and data extraction performance. Sharding Split Tables into a set of horizontal partitions or shards. This improves scalability. For any further detail: https://cloud.google.com/solutions/hybrid-and-multi-cloud-patterns-and-practices

Materialized View

Explanation:-Cache-Aside Load data on demand into a cache from a data store. This can improve performance and also helps to maintain consistency between data held in the cache and data in the underlying data store. If an application updates information, it can follow the write-through strategy by making the modification to the data store, and by invalidating the corresponding item in the cache. Materialized View Prepopulated views over the data in one or more tables when the data isn't ideally structured for required query operations. This can improve querying and data extraction performance. Sharding Split Tables into a set of horizontal partitions or shards. This improves scalability. For any further detail: https://cloud.google.com/solutions/hybrid-and-multi-cloud-patterns-and-practices

Sharding

Explanation:-Cache-Aside Load data on demand into a cache from a data store. This can improve performance and also helps to maintain consistency between data held in the cache and data in the underlying data store. If an application updates information, it can follow the write-through strategy by making the modification to the data store, and by invalidating the corresponding item in the cache. Materialized View Prepopulated views over the data in one or more tables when the data isn't ideally structured for required query operations. This can improve querying and data extraction performance. Sharding Split Tables into a set of horizontal partitions or shards. This improves scalability. For any further detail: https://cloud.google.com/solutions/hybrid-and-multi-cloud-patterns-and-practices

Queue-Based Load Leveling

Q48)

A company asked you to plan a systems migration and a new technological architecture for the Cloud.

You have been asked to indicate a series of patterns that can allow you to obtain greater efficiency in Availability.

Which of these options do you suggest (pick 2)?

Throttling

Explanation:-Throttling Control the consumption of resources used by an instance or a service; it allows applications to use resources only up to a limit, and then throttle them when this limit is reached. When usage exceeds the threshold, it can throttle requests from one or more users. This will enable the system to continue functioning and meet any service level agreements (SLAs) that are in place. Queue-Based Load Leveling Use a queue that acts as a buffer between a task and a service it invokes in order to smooth intermittent heavy loads that can cause the service to fail or the task to time out. The task and the service run asynchronously. For any further detail: https://cloud.google.com/solutions/hybrid-and-multi-cloud-patterns-and-practices

Cache-Aside

Materialized View

Sharding

Queue-Based Load Leveling

Explanation:-Throttling Control the consumption of resources used by an instance or a service; it allows applications to use resources only up to a limit, and then throttle them when this limit is reached. When usage exceeds the threshold, it can throttle requests from one or more users. This will enable the system to continue functioning and meet any service level agreements (SLAs) that are in place. Queue-Based Load Leveling Use a queue that acts as a buffer between a task and a service it invokes in order to smooth intermittent heavy loads that can cause the service to fail or the task to time out. The task and the service run asynchronously. For any further detail: https://cloud.google.com/solutions/hybrid-and-multi-cloud-patterns-and-practices

---

Q49) You have this SQL Statement in Bigquery: SELECT ANY_VALUE(fruit) as any_value FROM UNNEST(["apple", "banana", "pear"]) as fruit; What is the result of this query?

a table with 3 rows

apple

banana

pear

one of the values, randomly

Explanation:-The interesting elements in this query are: UNNEST: takes an ARRAY and returns a table, UNNEST may be used outside of the FROM clause with the IN operator Analytic Functions: a function that computes aggregate values over a group of rows. ANY_VALUE function that returns a random value from the input or NULL if there are zero input rows So, from the table created on the fly from an array, you may pick a casual value each time you execute the query. A more complex statement is: SELECT fruit, ANY_VALUE(fruit) OVER (ORDER BY LENGTH(fruit) ROWS BETWEEN 1 PRECEDING AND CURRENT ROW) AS any_value FROM UNNEST(["apple", "banana", "pear"]) as fruit; In the reference you will find a complete explanation. For any further detail: https://cloud.google.com/bigquery/docs/reference/standard-sql/analytic-function-concepts https://cloud.google.com/bigquery/docs/reference/standard-sql/query-syntax https://cloud.google.com/bigquery/docs/reference/standard-sql/aggregate_functions

---

Q50) What is the difference between Blue/green deployments, Traffic-splitting deployments, Rolling deployments and Canary deployments (picks 2)?

✅ Canary is a Rolling and Traffic-splitting deployment

Explanation:-Rolling deployments is a general technique that incrementally replaces old software with the new one. It is designed to update your workloads without downtime. Blue-green deployment lets have two production environments, as identical as possible, one old and one updated. The blue is live and you perform your final stage of testing in the green environment. Once the software is working in the green environment, you switch: the blue one is now idle and the green totally active. Canary is used to deploy in production new software versions by gradually rolling out the change to a small subgroup of users, before rolling it out to the entire platform/infrastructure and making it available to everybody. Canary deployment is like blue-green, but instead of switching from blue to green in one step, you use a phased approach. Traffic-splitting means that you have different environments and you divide the traffic among them. Any of these deployments lets you roll back to the previous stage. So, Canary is a Rolling and Traffic-splitting deployment because it is active on both versions. Blue-green is not a Traffic-splitting deployment because it is active on only 1 version. Managed deployment lets you ever roll back. For any further detail: https://cloud.google.com/solutions/continuous-delivery/ https://martinfowler.com/bliki/BlueGreenDeployment.html https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps https://cloud.google.com/appengine/docs/admin-api/migrating-splitting-traffic

⚫ Blue-green is a Rolling and Traffic-splitting deployment

✅ Canary and Blue-green let you roll back to the previous stage

Explanation:-Rolling deployments is a general technique that incrementally replaces old software with the new one. It is designed to update your workloads without downtime. Blue-green deployment lets have two production environments, as identical as possible, one old and one updated. The blue is live and you perform your final stage of testing in the green environment. Once the software is working in the green environment, you switch: the blue one is now idle and the green totally active. Canary is used to deploy in production new software versions by gradually rolling out the change to a small subgroup of users, before rolling it out to the entire platform/infrastructure and making it available to everybody. Canary deployment is like blue-green, but instead of switching from blue to green in one step, you use a phased approach. Traffic-splitting means that you have different environments and you divide the traffic among them. Any of these deployments lets you roll back to the previous stage. So, Canary is a Rolling and Traffic-splitting deployment because it is active on both versions. Blue-green is not a Traffic-splitting deployment because it is active on only 1 version. Managed deployment lets you ever roll back. For any further detail: https://cloud.google.com/solutions/continuous-delivery/ https://martinfowler.com/bliki/BlueGreenDeployment.html https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps https://cloud.google.com/appengine/docs/admin-api/migrating-splitting-traffic

⚫ Only Canary lets you roll back to the previous stage