

Google Cloud Digital Leader Practice Exam (Sample)

Study Guide



Everything you need from our exam experts!

Copyright © 2025 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain from reliable sources accurate, complete, and timely information about this product.

SAMPLE

Questions

SAMPLE

- 1. Which type of cloud implementation offers on-demand availability of computing and infrastructure resources?**
 - A. On-premises**
 - B. Private cloud**
 - C. Public cloud**
 - D. Hybrid cloud**
- 2. An organization is monitoring a service level indicator based on the volume of traffic that their application is receiving. Why is it important to monitor this indicator?**
 - A. The volume of traffic always affects the user experience.**
 - B. High traffic levels are frequently tied to degrading performance.**
 - C. Historical trends in traffic can be used for capacity planning.**
 - D. Low traffic volumes always indicate configuration issues.**
- 3. What characteristic describes software whose source code is publicly available and can be used, modified, and shared freely?**
 - A. Closed Source**
 - B. Open Standard**
 - C. Open Source**
 - D. Proprietary Software**
- 4. Which Google product or service is best suited for ingesting location data from onboard sensors in rental vehicles?**
 - A. Dataflow**
 - B. Pub/Sub**
 - C. Cloud SQL**
 - D. Cloud Storage**
- 5. Which service allows deployment and management of virtual machines in GCP?**
 - A. Google Kubernetes Engine**
 - B. Google App Engine**
 - C. Google Compute Engine**
 - D. Google Cloud Functions**

- 6. Which of the following factors can affect an application's performance due to delays in network communication?**
- A. Throughput**
 - B. Latency**
 - C. Bandwidth**
 - D. Redundancy**
- 7. An organization has a small app that sends a mobile notification to a customer whenever a new order is placed. They require a simple, event-driven, serverless service. Which service should they choose?**
- A. Cloud Functions**
 - B. Kubernetes Engine**
 - C. Cloud Run**
 - D. Compute Engine**
- 8. Which feature lets you set limits on the amount of resources that can be used by a project or user?**
- A. Quota policies**
 - B. Billing reports**
 - C. Budget alerts**
 - D. Committed use discounts**
- 9. What is the purpose of the Google Cloud Billing Report?**
- A. To monitor service health in real time**
 - B. To provide insights into the costs and usage of GCP services**
 - C. To forecast future expenses**
 - D. To track project milestones and deadlines**
- 10. Which Google Cloud product is best suited for structured or semi-structured data with a transactional workload and SQL capabilities?**
- A. Cloud Firestore**
 - B. Cloud Spanner**
 - C. BigQuery**
 - D. Cloud Bigtable**

Answers

SAMPLE

1. C
2. C
3. C
4. B
5. C
6. B
7. A
8. A
9. B
10. B

SAMPLE

Explanations

SAMPLE

1. Which type of cloud implementation offers on-demand availability of computing and infrastructure resources?

- A. On-premises**
- B. Private cloud**
- C. Public cloud**
- D. Hybrid cloud**

The public cloud is the type of cloud implementation that offers on-demand availability of computing and infrastructure resources. Public clouds are owned and operated by third-party cloud service providers who deliver their computing resources, such as servers and storage, over the internet. This environment allows users to scale resources up or down based on demand without the need to invest in or maintain physical infrastructure. In a public cloud, resources are shared among multiple clients, making it a cost-effective solution for individuals and businesses that require flexibility and can benefit from the economies of scale afforded by shared resources. Users can access applications and services like computing power, storage, and networking whenever they need them, utilizing a pay-as-you-go model that aligns with usage. The other types of cloud implementation do not inherently provide the same level of on-demand resource availability. On-premises setups require organizations to manage and maintain their physical infrastructure, limiting flexibility. A private cloud offers dedicated resources for a single organization, which can result in less flexibility and higher costs compared to public models. Hybrid cloud models blend public and private clouds, but the on-demand features are primarily driven by the public cloud component.

2. An organization is monitoring a service level indicator based on the volume of traffic that their application is receiving. Why is it important to monitor this indicator?

- A. The volume of traffic always affects the user experience.**
- B. High traffic levels are frequently tied to degrading performance.**
- C. Historical trends in traffic can be used for capacity planning.**
- D. Low traffic volumes always indicate configuration issues.**

Monitoring the volume of traffic that an application receives is crucial for several reasons, particularly regarding historical trends and capacity planning. By tracking the volume of traffic over time, organizations can identify patterns and fluctuations in user behavior, which allows them to forecast future demand. This information is vital for making informed decisions about resource allocation, scaling infrastructure, and ensuring that the application can handle peak loads without degrading performance. Monitoring traffic volume helps organizations proactively address potential issues before they impact the user experience. Planning based on historical traffic data can lead to optimized performance, reduced downtime, and an overall smoother operation of the application as it adapts to changing user needs. By anticipating traffic spikes or drops, organizations can allocate resources more effectively, ensuring that they meet user expectations and maintain service quality. In contrast, while high traffic levels can lead to performance degradation and low traffic levels may indicate configuration issues, these are not always the case and can be influenced by various external factors. Hence, the ability to analyze historical trends is a more strategic approach to managing application performance and reliability.

3. What characteristic describes software whose source code is publicly available and can be used, modified, and shared freely?

A. Closed Source

B. Open Standard

C. Open Source

D. Proprietary Software

The characteristic that describes software whose source code is publicly available and can be used, modified, and shared freely is known as open source. Open source software promotes a collaborative approach to software development, allowing individuals and organizations to access the source code, enhance it, and distribute their modifications. This model encourages innovation and transparency, as users can directly contribute to the software's evolution, ensuring that it evolves in ways that meet their needs. Open source software often comes with licenses that dictate how the software can be used and shared, but fundamentally, the key aspect is the freedom to inspect and modify the code. This creates an inclusive environment where developers and users can work together to improve the software for a variety of purposes, benefiting a wider community rather than being restricted to a single company or entity. In contrast, closed source and proprietary software restrict access to source code, limiting users' ability to modify or share the software. Open standards, while promoting interoperability and compatibility between different software systems, do not necessarily imply that the source code is available for modification or use.

4. Which Google product or service is best suited for ingesting location data from onboard sensors in rental vehicles?

A. Dataflow

B. Pub/Sub

C. Cloud SQL

D. Cloud Storage

The best choice for ingesting location data from onboard sensors in rental vehicles is Pub/Sub. This service is designed for high-throughput, real-time messaging, making it ideal for scenarios where data is generated continuously from sensors. Pub/Sub allows for seamless ingestion of data streams, enabling the rental vehicle system to handle and process location updates efficiently as they occur. As data is generated from multiple vehicles, Pub/Sub can scale appropriately, allowing the system to manage large volumes of incoming messages without delays. This is essential for real-time applications that require quick responses based on dynamic location data. In contrast, while Dataflow is powerful for real-time data processing and could be used in a broader context to process the ingested data, it functions primarily as a processing engine that would work in conjunction with a messaging service like Pub/Sub. Cloud SQL is a relational database useful for structured data storage rather than real-time ingestion of streaming data. Cloud Storage, while suitable for storing large amounts of data, is not optimized for real-time data ingestion and processing, as it is primarily designed for object storage rather than streaming. Thus, Pub/Sub is the most effective choice for this use case due to its design for real-time data ingestion and processing scalability.

5. Which service allows deployment and management of virtual machines in GCP?

- A. Google Kubernetes Engine**
- B. Google App Engine**
- C. Google Compute Engine**
- D. Google Cloud Functions**

The correct response identifies Google Compute Engine as the service designed specifically for the deployment and management of virtual machines (VMs) within Google Cloud Platform (GCP). Google Compute Engine provides Infrastructure as a Service (IaaS), allowing users to create and run VMs on Google's infrastructure. This service offers customizable VMs with various machine types, operating systems, and storage options, making it highly versatile and scalable based on the user's compute needs. By leveraging Google Compute Engine, users can efficiently manage workloads, scale resources dynamically, and benefit from Google's security and networking features. This capability is essential for businesses looking to host applications, run data analytics, or perform any compute-intensive tasks that require a flexible and reliable environment. In contrast, other services mentioned do not focus on virtual machine management. Google Kubernetes Engine is geared toward container orchestration rather than direct VM management, being optimized for running containerized applications. Google App Engine is a platform-as-a-service (PaaS) option that enables developers to build scalable web applications without worrying about the underlying infrastructure, which means it does not involve direct control over VMs. Google Cloud Functions is designed for serverless computing, executing code in response to events without needing to manage VMs at all, emphasizing

6. Which of the following factors can affect an application's performance due to delays in network communication?

- A. Throughput**
- B. Latency**
- C. Bandwidth**
- D. Redundancy**

Latency is a key factor that directly influences an application's performance due to delays in network communication. It refers to the time it takes for data to travel from the source to the destination. High latency can lead to noticeable delays when users interact with applications, resulting in a poor user experience. This can be particularly impactful in scenarios such as online gaming, video conferencing, or any real-time applications where response times are critical. While throughput and bandwidth are related to the volume of data that can be transmitted over the network, they do not specifically denote the delay in communication. Thus, they are more concerned with the capacity and speed of data transfer rather than the inherent time delays involved in the process. Redundancy, on the other hand, pertains to backup components that ensure reliability and availability of services but does not directly affect the performance related to network delays. Understanding the concept of latency and how it impacts application performance is crucial for designing and optimizing networked applications, especially in cloud environments where network characteristics can significantly vary.

- 7. An organization has a small app that sends a mobile notification to a customer whenever a new order is placed. They require a simple, event-driven, serverless service. Which service should they choose?**

A. Cloud Functions

B. Kubernetes Engine

C. Cloud Run

D. Compute Engine

For an organization that requires a simple, event-driven, serverless service to send mobile notifications upon new orders, Cloud Functions is the ideal choice. This service is designed specifically for executing code in response to events, making it a great fit for the use case presented. It allows developers to run their code without managing the underlying infrastructure, thus streamlining the development process. Cloud Functions enables automatic scaling depending on the number of incoming events, ensuring that it can handle varying loads without needing any manual intervention. In this scenario, when a new order event occurs, Cloud Functions can be triggered to send a notification seamlessly. The other options, while reliable services in their own right, do not align as well with the need for a simple, serverless, event-driven architecture. Kubernetes Engine, for example, is more suited for deploying containerized applications where management of clusters is necessary, which adds complexity that isn't required for this scenario. Cloud Run serves containerized applications but involves somewhat more overhead in terms of container management. Compute Engine requires provisioning and managing virtual machines, making it the least ideal choice for a lightweight, event-driven notification service. Thus, Cloud Functions stands out as the best option for responding to specific events while remaining serverless.

- 8. Which feature lets you set limits on the amount of resources that can be used by a project or user?**

A. Quota policies

B. Billing reports

C. Budget alerts

D. Committed use discounts

The correct choice, quota policies, is designed specifically to manage and control the usage of resources within Google Cloud. Quota policies allow administrators to define limits on various resources such as APIs, storage, and compute instances per project or user. This feature is particularly important for ensuring that resource availability is maintained and that costs are managed effectively, thus preventing unexpected charges or resource overuse. Quota policies are valuable for organizations that need to enforce usage limitations or prioritize resources for particular projects or users. By setting these limits, organizations can better control their resource spending, avoid capacity issues, and ensure compliance with internal governance and budgetary requirements. In contrast, billing reports provide insights into past usage and expenditures rather than actively managing or limiting resource allocation. Budget alerts notify users when spending approaches or exceeds predefined budgets but do not impose usage restrictions. Committed use discounts are pricing options that provide cost savings when committing to using certain resources over a specified period but do not relate to resource limitations. Thus, quota policies are the foundational feature for managing and setting limits on resource usage effectively.

9. What is the purpose of the Google Cloud Billing Report?

- A. To monitor service health in real time
- B. To provide insights into the costs and usage of GCP services**
- C. To forecast future expenses
- D. To track project milestones and deadlines

The purpose of the Google Cloud Billing Report is to provide insights into the costs and usage of Google Cloud Platform (GCP) services. This tool enables organizations to analyze their expenditures, track spending trends, and understand how different services contribute to overall costs. By visualizing billing data, users can make informed decisions regarding resource allocation, budget adjustments, and cost management strategies. Gaining detailed insights into costs can help organizations identify areas where they could optimize spending and better manage their budgets. The billing report offers features like filtering and grouping data by different dimensions, such as projects, services, and labels, allowing users to tailor their analysis to meet specific financial needs. Understanding these cost insights is crucial for effective financial planning and resource management within GCP. This analysis supports users in making informed decisions about cloud service usage, helping to ensure that the organization remains within budget while meeting its operational needs.

10. Which Google Cloud product is best suited for structured or semi-structured data with a transactional workload and SQL capabilities?

- A. Cloud Firestore
- B. Cloud Spanner**
- C. BigQuery
- D. Cloud Bigtable

Cloud Spanner is particularly well-suited for structured or semi-structured data, especially in environments where transactional workloads are a key requirement. It combines the benefits of traditional relational databases with horizontal scalability, which allows it to handle large amounts of data and high transaction volumes while maintaining consistency. One of Cloud Spanner's standout features is its support for SQL, enabling users to perform complex queries on their data just as they would with other relational database systems. This makes it an excellent choice for applications that require strong consistency, ACID transactions, and global distribution. In contrast, options like Cloud Firestore and Cloud Bigtable are more optimized for specific use cases. Firestore is tailored for real-time data syncing and document-oriented applications, making it less ideal for transactional SQL workloads. Cloud Bigtable excels in handling high-throughput analytical workloads and is designed for large volumes of unstructured or semi-structured data rather than transactional operations. BigQuery, while powerful for large-scale analytics and structured data processing, is not designed for transactional workloads and does not support transactional SQL operations in the same way that Cloud Spanner does. Together, these differences highlight why Cloud Spanner is the best choice for scenarios that require SQL capabilities, structured data management, and reliable transactional support.