

# 111 GCP interview questions to hire an expert

## Questions

1. What is Google Cloud Platform in simple terms?
2. Can you name a few GCP services you know about?
3. What is the difference between Compute Engine and App Engine?
4. Explain what a GCP project is and why it's important.
5. What is Google Cloud Storage, and when would you use it?
6. How do you access GCP services? Name a couple of ways.
7. What is the purpose of Identity and Access Management (IAM) in GCP?
8. What's the difference between a region and a zone in GCP?
9. Have you ever used the Google Cloud Shell? What is it for?
10. What is a virtual machine, and how does it relate to Compute Engine?
11. What is the purpose of VPC (Virtual Private Cloud) in GCP?
12. How can you monitor the performance of your applications in GCP?
13. What are some ways to keep your data safe in Google Cloud?
14. Describe a situation where you might use Cloud Functions.
15. What are some advantages of using a cloud platform like GCP?
16. Explain what a container is and if you know of any container service in GCP.
17. What's the basic idea behind autoscaling in Compute Engine?
18. If a service goes down in one zone, what can you do to ensure high availability?
19. What is BigQuery and when would you use it?
20. What are some of the free resources available to new GCP users?
21. How can you estimate the cost of running a workload on GCP?
22. What are some key differences between SQL and NoSQL databases, and which does GCP offer?
23. What is a service account in GCP, and why is it useful?
24. Explain the difference between preemptible and non-preemptible VMs in GCP and when would you choose one over the other?
25. How do you handle rolling updates and rollbacks for applications running on Google Kubernetes Engine (GKE)?
26. Describe the purpose of Cloud IAM roles and permissions. How do you grant least privilege access to a service account?
27. What are the differences between Cloud Storage Nearline, Coldline, and Archive storage classes, and how do you decide which one to use?
28. How can you monitor the performance of your applications running on GCP, and what metrics are most important to track?
29. Explain how to set up a basic CI/CD pipeline using Cloud Build for a containerized application.
30. What is the purpose of VPC Service Controls, and how does it enhance security in GCP?
31. How can you automate infrastructure provisioning and management in GCP using Infrastructure as Code (IaC) tools like Terraform?
32. Describe the differences between Cloud SQL and Cloud Spanner, and when would you use each?
33. Explain how to implement a hybrid cloud environment connecting your on-premises infrastructure to GCP.
34. How do you configure and use Cloud Load Balancing for distributing traffic across multiple instances?
35. What are the different ways to authenticate to GCP services, and when should you use each method?
36. Explain the purpose of Kubernetes namespaces and how they help in organizing resources within a GKE cluster.
37. How can you optimize the cost of your GCP resources? Explain various cost management strategies.
38. Describe the process of setting up a VPN connection between your on-premises network and a GCP VPC.
39. What are the benefits of using managed instance groups (MIGs), and how do they improve application availability?
40. How do you handle secrets management in GCP, and what tools/services can you use to store and access secrets securely?
41. Explain the difference between Cloud Functions and Cloud Run, and when would you choose one over the other for serverless deployments?
42. How can you use Cloud Monitoring and Cloud Logging to troubleshoot issues with your applications in GCP?
43. Describe the purpose of Cloud CDN and how it can improve the performance of your web applications.
44. Explain how to use Cloud Dataproc for running Apache Spark and Apache Hadoop jobs in GCP.
45. How can you ensure data residency and compliance requirements are met when storing data in GCP?
46. What are the benefits of using gcloud CLI and how do you manage multiple GCP projects with it?
47. Explain how to implement a disaster recovery plan for your applications running in GCP.
48. How do you configure autoscaling for your applications in GKE based on resource utilization?
49. Describe the different types of network peering options available in GCP and their use cases.
50. Explain how to use Cloud Pub/Sub for building asynchronous messaging systems in GCP.
51. How do you secure your GKE cluster and protect it from unauthorized access and vulnerabilities?
52. What is the purpose of Cloud Armor and how can it protect your applications from DDoS attacks and other web threats?
53. How can you ensure data consistency across multiple regions in a globally distributed Cloud Spanner database?
54. Describe a scenario where you would choose Memorystore over Cloud SQL for caching data, and explain why.
55. Explain how you would implement a CI/CD pipeline for a serverless application using Cloud Build, Cloud Functions, and Cloud Deploy.
56. How do you monitor and troubleshoot performance bottlenecks in a microservices architecture running on GKE?
57. Explain the difference between preemptible VMs and standard VMs, and describe a use case for each.
58. How can you optimize the cost of running a large-scale data processing job on Dataproc?
59. Describe a strategy for implementing disaster recovery for a critical application running on GCP, including RTO and RPO considerations.
60. How would you secure a GKE cluster to meet specific compliance requirements, such as PCI DSS or HIPAA?
61. Explain how you can use Cloud IAM to implement fine-grained access control for different teams working on the same GCP project.
62. Describe the process of migrating a large on-premises database to Cloud SQL with minimal downtime.
63. How can you use Cloud CDN to improve the performance and availability of a website hosted on Compute Engine?
64. Explain how you would use Cloud Monitoring and Cloud Logging to detect and respond to security incidents in your GCP environment.
65. How do you manage and automate infrastructure as code using Terraform on GCP?
66. Describe how you would implement a hybrid cloud solution that connects your on-premises infrastructure to GCP.
67. How can you use Cloud Functions to build event-driven applications that respond to changes in Cloud Storage?
68. Explain how you would design a data lake on Cloud Storage for analytics and machine learning workloads.
69. How do you ensure data privacy and compliance when processing sensitive data in BigQuery?
70. Describe how you would use Cloud Composer to orchestrate complex data pipelines.
71. How can you use Cloud Run to deploy and scale containerized applications without managing servers?
72. Explain how you would implement a blue/green deployment strategy for a critical application running on GKE.
73. How can you leverage Vertex AI for building and deploying machine learning models on GCP?
74. Design a solution for near-real-time data ingestion and analysis using Pub/Sub, Dataflow, and BigQuery. Discuss scalability and fault tolerance.
75. Describe the differences between Cloud SQL, Cloud Spanner, and Datastore. In which scenarios would you choose each one, and why?
76. How would you implement a multi-region, active-active architecture for a web application hosted on Compute Engine? Consider data replication, load balancing, and failover strategies.
77. Explain how you can use Cloud Armor to protect your web applications from common web exploits and attacks, such as SQL injection and cross-site scripting.
78. Outline a strategy for migrating a large-scale, stateful application from AWS to GCP. What are the key considerations and potential challenges?
79. Describe how you would use Cloud KMS to manage encryption keys for data stored in Cloud Storage and BigQuery.
80. Explain the purpose of service accounts in GCP and how you would use them to grant permissions to applications running on Compute Engine or GKE.
81. How would you design a cost-effective solution for storing and archiving large volumes of infrequently accessed data on GCP?
82. Considering a hybrid cloud setup with on-premises data center and GCP, how to effectively manage the networking and security policies?
83. How do you design a highly available and scalable data pipeline using Dataflow, considering various failure scenarios and data consistency requirements?
84. Explain how you would implement a multi-region deployment strategy for a critical application on GKE, detailing the failover mechanisms and data replication techniques?
85. Describe your approach to securing a hybrid cloud environment that connects your on-premises infrastructure to GCP, focusing on identity management and network security?
86. How would you optimize the cost of a large-scale BigQuery deployment while maintaining query performance and data availability?
87. Design a solution for real-time fraud detection using Pub/Sub, Dataflow, and BigQuery, considering the challenges of low latency and high accuracy.
88. Explain how you would implement a blue/green deployment strategy for a microservices application on GKE, ensuring minimal downtime and seamless rollback capabilities?
89. Describe your experience with implementing infrastructure as code (IaC) using Terraform in GCP, focusing on modularity, reusability, and version control?
90. How would you troubleshoot performance bottlenecks in a complex application running on App Engine, considering various factors such as database queries, network latency, and code efficiency?
91. Design a disaster recovery plan for a critical application running on Compute Engine, considering various failure scenarios and recovery time objectives (RTO)?
92. Explain how you would implement a data governance framework for a data lake in Cloud Storage, focusing on data quality, metadata management, and access control?
93. Describe your approach to implementing a CI/CD pipeline for a serverless application using Cloud Functions and Cloud Build, emphasizing automated testing and deployment?
94. How would you optimize the performance of a machine learning model deployed using Vertex AI, considering factors such as model size, inference latency, and resource utilization?
95. Design a solution for ingesting and processing streaming data from IoT devices using Pub/Sub, Dataflow, and Bigtable, considering the challenges of high volume and low latency?
96. Explain how you would implement a multi-tenancy architecture for a SaaS application on GCP, focusing on data isolation, resource allocation, and security?
97. Describe your experience with implementing security best practices for Kubernetes deployments on GKE, focusing on network policies, pod security policies, and RBAC?
98. How would you design a solution for migrating a large on-premises database to Cloud SQL, considering the challenges of data transfer, schema migration, and application downtime?
99. Explain your approach to monitoring and alerting for a distributed application running on GCP, focusing on key performance indicators (KPIs) and root cause analysis?
100. Describe how you would implement a data encryption strategy for data at rest and in transit across various GCP services, focusing on key management and compliance requirements?
101. How do you handle secrets management in a GCP environment, especially when dealing with sensitive information in configuration files and code repositories?
102. Explain your approach to implementing a role-based access control (RBAC) strategy for a complex GCP environment with multiple teams and projects?
103. How would you design a cost-effective and scalable solution for storing and serving large media files (e.g., images, videos) using Cloud Storage and Cloud CDN?
104. Explain how you would implement a secure and reliable VPN connection between your on-premises network and a VPC in GCP, considering various security and performance factors?
105. Describe your experience with using service mesh technologies like Istio on GKE, focusing on traffic management, security, and observability?
106. How would you design a solution for automatically scaling a web application running on Compute Engine based on real-time traffic patterns and resource utilization?
107. Explain how you would implement a data anonymization or pseudonymization strategy for sensitive data stored in BigQuery, complying with privacy regulations like GDPR?
108. Describe your approach to implementing a hybrid cloud storage solution that combines the benefits of on-premises storage with the scalability and cost-effectiveness of