

Q1. What is BigQuery and what is its primary value proposition within its cloud ecosystem?

Answer:

Q2. How does BigQuery handle identity and access management (IAM)? Explain roles, policies, and permissions.

Answer:

Q3. What networking primitives does BigQuery use (VPC, subnets, security groups, load balancers)?

Answer:

Q4. Explain the storage options available in BigQuery and when you would choose each.

Answer:

Q5. How does BigQuery handle autoscaling and how do you configure it for cost efficiency?

Answer:

Q6. What monitoring and observability tools does BigQuery provide (metrics, logs, tracing)?

Answer:

Q7. Explain how BigQuery handles high availability and disaster recovery.

Answer:

Q8. How do you manage secrets and configuration in BigQuery (secret managers, parameter stores)?

Answer:

Q9. What are the key cost drivers when using BigQuery and how do you optimize spend?

Answer:

Q10. How do you implement CI/CD pipelines that deploy to BigQuery?

Answer:
