

Q1. What is Amazon Redshift and what primary problem does it solve in the Data Platforms space?

Answer:

Q2. Explain the core architecture or components of Amazon Redshift and how they work together.

Answer:

Q3. What are the main alternatives to Amazon Redshift and when would you choose it over them?

Answer:

Q4. How do you install, configure, and get started with Amazon Redshift for a new project?

Answer:

Q5. What are the key configuration options or settings that affect Amazon Redshift's behavior in production?

Answer:

Q6. How does Amazon Redshift handle scalability, reliability, and high availability?

Answer:

Q7. What observability does Amazon Redshift provide: metrics, logs, traces, or health endpoints?

Answer:

Q8. What are the common performance bottlenecks in Amazon Redshift and how do you address them?

Answer:

Q9. What security best practices should you follow when deploying Amazon Redshift?

Answer:

Q10. What are the most common pitfalls or mistakes teams make when adopting Amazon Redshift?

Answer:
