

**Q1.** What is Qdrant and what machine learning or AI problems is it designed to solve?

Answer:

**Q2.** Explain the core abstractions or APIs in Qdrant (models, tensors, pipelines, agents).

Answer:

**Q3.** How does Qdrant handle training: defining a model, specifying a loss, and running optimization?

Answer:

**Q4.** What is Qdrant's approach to data ingestion, preprocessing, and batching?

Answer:

**Q5.** How do you evaluate model performance in Qdrant (metrics, validation sets, cross-validation)?

Answer:

**Q6.** How does Qdrant support model deployment and serving in a production environment?

Answer:

**Q7.** What hardware acceleration does Qdrant support (GPU, TPU) and how do you configure it?

Answer:

**Q8.** How does Qdrant handle distributed or multi-GPU training?

Answer:

**Q9.** What are common debugging techniques for model training issues in Qdrant?

Answer:

**Q10.** How do you version and experiment-track models in a Qdrant-based ML workflow?

Answer: