

Data Intensive Computing - Review Questions 1

Deadline: September 9, 2023

1. Explain how a file region can be left in an **undefined** state in GFS?

2. Briefly explain the read operation in GFS? What's the role of the primary node?

3. Using one example show that in the CAP theorem, if we have Consistency and Partition Tolerance, we cannot provide Availability at the same time.

4. Explain how the consistent hashing works.

5. Explain the finding process of tablets in BigTable.