

General Instruction: After the seventh lecture of big data course, you should be able to answer the following questions.

1. What are four key motivations of GFS in Google workloads and technological environment?
2. What are application scenarios where GFS/HDFS are **NOT** a good fit for?
3. What are the main differences between a traditional relational DBMS and a big data management system?
4. What is the basic data model of Bigtable? Does it conform to the first normal form of relational databases?

Hint: A relation is in first normal form if and only if the domain of each attribute contains only atomic (indivisible) values, and the value of each attribute contains only a single value from that domain.

5. What are three major components of Bigtable? What are their main functions?
6. What are five types of NoSQL databases?
7. Considering the following table, show three ways to store the data in row-oriented databases, column-oriented databases and key-value databases, respectively.

RowId	Student Id	Last-name	First-name	Age
001	10	Smith	Joe	25
002	12	Jones	Mary	21
003	11	Johnson	Cathy	34
004	22	Jones	Bob	42