


Name:			
Enrolment No:			
<b>UPES</b> <b>End Semester Examination, May 2024</b>			
<b>Course: Big Data Storage</b> <b>Program: B.Tech. (CSE+BIG DATA-H/NH)</b> <b>Course Code: CSBD2001</b>		<b>Semester: IV</b> <b>Time: 03 hrs.</b> <b>Max. Marks: 100</b>	
<b>Instructions:</b> Attempt all questions			
<b>SECTION A</b> <b>(5Qx4M=20Marks)</b>			
S. No.		Marks	CO
Q 1	List out different types of storage Media.	4	CO1
Q 2	Define Client-Server Architecture and Symmetric Architecture of DFS.	4	CO1
Q 3	Define the features and use cases of Columnar Databases.	2+2	CO2
Q 4	In HBase, assume table 'Student' with column family "Academic_Details". Write HBase shell command for: a) creating table 'Student' b) storing a column 'cgpa' with value '6.7' and key as 'R21' c) update value of row where key is 'R21' and column is 'cgpa' to '7.1' d) dropping table 'Student'	4	CO3
Q 5	Why is block size a crucial design consideration in HDFS?	4	CO4
<b>SECTION B</b> <b>(4Qx10M= 40 Marks)</b>			
Q 6	Explain the CAP theorem and its three properties. How do these properties interrelate? In what scenarios might eventual consistency be an acceptable trade-off for a distributed system?	6+4	CO1
Q 7	Imagine you're building a new social media platform. Would you choose a traditional database or a NoSQL database? Why? Explain the pros and cons of each option.	10	CO1
Q 8	NFS operates in a client-server model. Describe the architecture and functionalities handled by the NFS client and the NFS server, respectively. How do they interact with each other to facilitate remote file access?	5+5	CO2

Q 9	<p>Write short notes on <b>any 2</b> topics:</p> <ul style="list-style-type: none"> <li>a) File Caching in DFS</li> <li>b) Superblock</li> <li>c) Quotas &amp; Access Control Lists</li> <li>d) VMware VMFS</li> </ul>	5+5	CO2
<b>SECTION-C</b> <b>(2Qx20M=40 Marks)</b>			
Q 10	<p>Discuss various factors for maintaining robustness and achieving fault tolerance in HDFS.</p> <p style="text-align: center;"><b>OR</b></p> <p>Break down the steps involved in HDFS client-server interaction for both read and write operations.</p>	20	CO3
Q 11	<p>In what scenarios might sharding not be an effective scaling strategy? Compare different sharding strategies in detail.</p>	8+12=20	CO4