Course Code	Course Title	L	T	P	C
BCSE406L	NoSQL Databases	3	0	0	3
Pre-requisite	NIL	Syllabus version			
		1.0			

Course Objectives

1. To recognize the emergence, requirements and benefits of a NoSQL database.
2. To compare NoSQL databases with each other and relational systems.
3. To create wide-column, document, key-value, graph and object-oriented databases, add content, and run queries.

Course Outcomes

1. Define, compare and use the four types of NoSQL Databases.
2. Explain key value databases and apply queries on those databases.
3. Explain the detailed architecture, define objects, load data, query data and performance tune Document-oriented NoSQL databases.
4. Demonstrate an understanding of the detailed architecture, define objects, load data, query data and performance tune Column-oriented and Graph oriented NoSQL databases.
5. Evaluate NoSQL database development tools and programming languages.

Module:1	Introduction	5 hours

The Relational Database Revolution with SQL - Design Limitations - Impedance Mismatch - Schema Evolution - Horizontal Scalability- Motivations for Not Just/No SQL (NoSQL) Databases - Data Management with Distributed Databases – SQL Versus NoSQL Databases - The CAP Theorem - ACID and BASE - Types of Eventual Consistency - Types of NoSQL Databases

Module:2	Key-Value Databases	4 hours

From Arrays to Key-Value Databases - Essential Features of Key-Value Databases - Properties of keys - Characteristics of values - Data Modeling Terms for Key- Value Databases - Key - Architecture and Implementation terms - Designing Structured Values - Limitations of Key-Value Databases - Design Patterns for Key-Value Databases - Redis database - Queries - Case Study: Key-Value Databases for Mobile Application Configuration.

Module:3	Document Databases	6 hours

Introduction to Document Databases – Mongo DB - Collections - Basic operations on collections (CRUD) - Find operation - Sorting - Limiting - Aggregate operations: Aggregation pipeline - Operators - Combining aggregate operators

Module:4	Designing Document Databases	7 hours

Partitioning - Types of Partitions - Vertical Partitioning - Horizontal Partitioning or Sharding - Separating Data with Shard Keys – Replication - Distributing Data with a Partitioning Algorithm - Data Modelling and Query Processing - Normalization – De-Normalization - Query Processor - Indexing - Distributed Consistency - Joins

Module:5	Column Family Databases	6 hours

Introduction to Column Family Databases - Google Big Table - Differences and Similarities to Key-Value and Document Databases - Architectures Used in Column Family Databases - Cassandra Architecture: Peer-to-Peer - Commit Log - Bloom Filter - Consistency Level - Processes and Protocols - Replication - Anti-Entropy -

Gossip Protocol - Hinted Handoff - Handling of deletion - When to Use Column Family		
Module:6	**Designing Column Family Databases**	**7 hours**
Column Family Database Terminology – Key space - Row Key - Column - Column Families – Cassandra - CQL Queries - Primary Key and Clustering Key - CRUD operations - Cluster - Partition - Replication - Consistency levels - Guidelines for Designing Tables - Indexing - Primary and Secondary Index		
Module:7	**Graph Databases**	**8 hours**
Introduction to Graph Databases - What Is a Graph? - Graphs and Network Modeling - Advantages of Graph Database - Elements of Graphs - Operations on Graphs - Properties of Graphs and Nodes - Types of Graphs - Graph Design - Neo4J - Cypher Query Language: Creating, Removing and Querying Nodes and Relations - Basic Graph Traversal - Finding Path and Distance between nodes - Gremlin: Query by Graph Traversal - Using NoSQL and Relational Databases Together		
Module:8	**Recent Trends**	**2 hours**
Guest lectures from Industry and, Research and Development Organizations		
	Total Lecture hours:	**45 hours**

Text Book(s)	
1.	Dan Sullivan, NoSQL for Mere Mortals, Addison-Wesley Professional, 2015
Reference Books	
1.	Adam Fowler, NoSQL For Dummies, For Dummies, 1st edition, 2015
2.	Gerardus Blokdyk, NoSQL Databases A Complete Guide, 5STARCooks, 2021
3.	Pramod J. Sadalage and Martin Fowler, NoSQL Distilled, Addison-Wesley, 1st edition, 2012
4.	Guy Harrison, "Next Generation database: NoSQL New SQL and Big Data", Apress, 1st Edition, 2015.

Mode of Evaluation: CAT / written assignment / Quiz / FAT

Recommended by Board of Studies	12-05-2023		
Approved by Academic Council	No. 70	Date	24-06-2023