Kadi Sarva Vishwavidyalaya, Gandhinagar MCA Semester II MCA-26 (B) : Next Generation Databases

Rationale: The primary objective of this course is to provide in-depth knowledge of the next generation databases and future database technologies from SQL to NoSQL to NewSQL, different databases like document, graph, columnar databases etc.

Prerequisite: Knowledge of DBMS, ADBM, SQL & PL/SQL is desirable.

Learning Outcomes:

Students will learn seven concepts like database revolution, Document Databases, Graph Databases, Column Databases, In-Memory Databases, Object Databases, and Databases of Future. At the end of this course, they will enhance their conceptual and analytical understanding as under:

- In database revolution they will understand the first, second and third revolution of databases from file base databases to SQL to NoSQL to NewSQL.
- In document databases, they will be able to understand the document database overview, NoSQL databases, overview of MongoDB and CouchDB.
- In Graph and Column databases, they will be able to understand the graph and column databases overview and overview of Neo4j, Gremlin, Sybase IQ, C-Store, and Vertica.
- In Memory and Object databases, they will be able to understand features of in-memory databases and object databases overview and overview of TimesTen , Redis , SAP HANA, VoltDB , Oracle 12c, SPARK architectures.
- In last section of the course they will be able to make them understands about future database technologies like storage, block chain and quantum computing.
- All these concepts are important to build their career as Data Architecture/Data Engineers, Data Scientist, Data Analyst and DBAs.

Teaching and Evaluation Scheme: The objective of evaluation is to evaluate the students throughout the semester for better performance. Students are evaluated on the basis of continuous evaluation system both in theory and practical classes based on various parameters like term work, class participation, practical and theory assignments, presentation, class test, Regular Attendance, etc.

Sub Total Credit	Teaching scheme		Examination scheme				
	(per week)		MID	CEC	External		Total
	Th	Pr	Th	Th	Th.	Pr.	Marks
3	3	-	25	25	50	-	100

Course content:

Unit 1 Database Revolution

Three Database Revolutions, Early Database Systems-The First Database Revolution, The Second Database Revolution- Relational theory, Transaction Models, The First Relational Databases, Clientserver Computing, Object-oriented Programming and the OODBMS, The Relational Plateau, The Third Database Revolution, Google and Hadoop, The Rest of the Web, Cloud Computing, Document Database, The "NewSQL", The Nonrelational Explosion. Google, Big Data, and Hadoop, The Big Data Revolution- Cloud, Mobile, Social, and Big Data, Google: Pioneer of Big Data, Google Hardware, The Google Software Stack, More about MapReduce, Hadoop: Open-Source Google Stack -Hadoop's Origins, The Power of Hadoop, Hadoop's Architecture, HBase, Hive, Pig, The Hadoop Ecosystem.

Unit 2 Document Databases

is a document database, NoSQL databases, Why choose NoSQL?, Performance overview of different databases, Why a document store, How does it work, Data storage ,Data querying and the map/reduce paradigm ,Inserting and Modifying, ACID, The different solutions -Open source solution Proprietary solution. Examples - CouchDB, Why CouchDB, The storage, concurrency, Managing the database, Querying the database, Specificity of Couch DB.

Examples - MongoDB, Why MongoDB?, The storage, concurrency, Managing the database, Querying the database, Specificity of Mongo DB.

Unit 3 Graph Databases & Column Databases

What is a Graph?, RDBMS Patterns for Graphs, RDF and SPARQL, Property Graphs and Neo4j, Gremlin, Graph Database Internals, Graph Compute Engines. What is Column Databases, Why it used? The Columnar Alternative - Columnar Compression, Columnar Write Penalty, Sybase IQ, C-Store, and Vertica, Column Database Architectures -Projections, Columnar Technology in Other Databases.

Unit 4 In-Memory Databases & Object Databases

What is In-Memory Databases?, The End of Disk? -Solid State Disk , The Economics of Disk ,SSD-Enabled Databases, In-Memory Databases-Examples : TimesTen, Redis, SAP HANA, VoltDB, Oracle 12c "in-Memory Database", Berkeley Analytics Data Stack and Spark, Spark Architecture. Overview of object databases, Object Oriented Database, Object Relational Database, mapping of object relational mapping and standards of ODBMS, Examples -ObjectDB, ObjectStore, ObjectivityDB, db4o and GemStone features and advantages.

Unit 5 Databases of Future

The revolution revisited, counterrevolutionaries-have we come full circle?, can we have it all?consistency models, schema, database languages, storage, a vision for a converged database, other convergent databases, Disruptive database technologies-storage technologies, Blockchain-What it is?, Understanding Technologies, When it is used?, Quantum computing-Quantum Transaction, Quantum Search, Quantum Query Language.

Text Books:

- 1. Next Generation Databases NoSQL, NewSQL and Bigdta, Guy Harrison, Apress.
- 2. CouchDB, Document oriented Databases, Alain Issa, François Schiltz, ULB
- 3. Document stores and MongoDB, Kaïs Albichari, Tanguy d'Hose, ULB
- 4. MongoDB Architecture Guide, MongoDB university, white paper
- 5. Graph Databases-neo4j, Ian Robinson, Jim Webber & Emil Eifrem, 2nd edition .Oreilly

[20%]

[20%]

[20%]

[20%] What

[20%]

- 6. The Design and Implementation of Modern Column-Oriented Database Systems, Daniel Abadi, Foundations and Trends in Databases Vol. 5, No. 3 (2012) 197–280, 2013
- 7. http://justinlevandoski.org/papers/fnt-mmdb.pdf
- 8. Oracle® Database, Database In-Memory Guide, 12c Release, Lance Ashdown, Oracle Press
- 9. Fundamentals of Object Databases: Object-Oriented and Object-Relational Design,SuzanneW. Dietrich and Susan D. Urban, Morgan & cLaypool publishers
- 10. Blockchain basics, Technical Introduction in 25 Steps, Daniel Drescher, Apress

Web References:

- 1. couchdb.apache.org
- 2. CouchDB: The Definitive Guide:guide.couchdb.org/
- 3. Amazon DynamoDB vs. CouchDB vs. MongoDB Comparison
- 4. https://db-engines.com/en/system/Amazon+DynamoDB%3BCouchDB%3BMongoDB
- 5. https://university.mongodb.com/
- 6. https://pdfs.semanticscholar.org/f511/7084ca43e888fb3e17ab0f0e684cced0f8fd.pdf
- 7. http://asiandatascience.com/wp-content/uploads/2017/12/Definitive-Guide-Graph-Databases-for-RDBMS-Developer.pdf
- 8. http://www.durusau.net/localcopy/Graph-Modeling-Do.s-and-Don.ts.pdf
- 9. https://www.jugdo.de/wp-content/uploads/2014/01/intro_to_neo4j_jugdo.pdf
- 10. http://nms.csail.mit.edu/~stavros/pubs/tutorial2009-column_stores.pdf