CLOUD COMPUTING AND DISTRIBUTED SYSTEMS

Subject Code: CS35212CS	**Subject Type:** DC	**Evaluation Scheme:** MSE (30), TA (20), ESE (50)	**L-T-P:** 2-0-2 (3)

Pre-Requisites: Operating Systems, Computer Networks, Data Structures and Algorithms, Database Systems, Programming (Python/Java), Basic Knowledge of Virtualization and Parallel Computing

Course Objective: To provide a comprehensive understanding of the principles, architectures, and technologies underlying cloud computing and distributed systems. To enable students to design, analyze, and deploy scalable, fault-tolerant, and secure distributed applications in modern cloud environments.

Course Outcomes: At the end of the course, the student will be able to

CO	Course Outcomes
CO-1	Analyze the fundamental principles and theoretical models of distributed systems and cloud computing
CO-2	Design scalable and fault-tolerant distributed architectures using modern cloud paradigms
CO-3	Implement distributed computing frameworks and cloud-native applications
CO-4	Evaluate performance, security, and reliability of cloud-based systems in real-world scenarios

Syllabus:

Unit-1: Foundations of Distributed Systems

Models of distributed systems: synchronous vs asynchronous systems, system architectures (client-server, peer-to-peer, microservices), Communication paradigms: RPC, RMI, message passing, publish-subscribe systems, Time and ordering: logical clocks, vector clocks, causality. Consistency models: strong consistency, eventual consistency, CAP theorem, Distributed coordination: consensus (Paxos, Raft), leader election, distributed mutual exclusion.

Unit-2: Cloud Computing Architecture and Virtualization

Cloud computing fundamentals: service models (IaaS, PaaS, SaaS), deployment models (public, private, hybrid), Virtualization technologies: hypervisors, containers, container orchestration (Docker, Kubernetes), Resource management and scheduling in cloud environment, Data center design and energy-efficient computing, Serverless computing and Function-as-a-Service (FaaS).

Unit-3: Distributed Data Management and Big Data Systems

Distributed file systems: GFS, HDFS, Distributed databases: NoSQL systems (key-value, column-family, document stores), Data partitioning, replication, and sharding strategies, Consistency, availability, and partition tolerance trade-offs, Big data processing frameworks: MapReduce, Apache Spark, Stream processing systems and real-time analytics.

Unit-4: Scalability, Fault Tolerance, and Security in Cloud Systems

Scalability techniques: load balancing, auto-scaling, elastic resource provisioning, Fault tolerance: replication, checkpointing, failure detection, recovery mechanisms, Monitoring and observability in distributed systems, Cloud security: identity and access management (IAM), data security, encryption, zero-trust architecture, Privacy, compliance, and governance in cloud systems, Emerging trends: edge computing, fog computing, blockchain integration with cloud.

Learning Resources:

Text Books:

1. George Coulouris, Jean Dollimore, Tim Kindberg, Distributed Systems: Concepts and Design, Pearson.
2. Kai Hwang, Geoffrey Fox, Jack Dongarra, Distributed and Cloud Computing, Morgan Kaufmann.

Reference Books:

1. Martin Kleppmann, Designing Data-Intensive Applications, O'Reilly.
2. Brendan Burns, Kubernetes: Up and Running, O'Reilly.
3. Tanenbaum & Van Steen, Distributed Systems, Pearson.
4. Rajkumar Buyya et al., Cloud Computing: Principles and Paradigms, Wiley.