Registration No.: | | | | | | | | | | |

Total Number of Pages: 02 **Course: B.Tech**

Sub_Code: PCAC2011

4th Semester Regular Examination: 2024-25

SUBJECT: CLOUD INFRASTRUCTURE AND APPLICATIONS

BRANCH(S): AEIE, CIVIL, CSE, CSEAI, CST, ECE, EEE, ELECTRICAL, ETC, MECH

Time: 3 Hours

Max Marks: 100

Q.Code: S272

Answer Question No.1 (Part-I) which is compulsory, any eight from Part-II and any two from Part-III.

The figures in the right hand margin indicate marks.

Part-I

Q1 **Answer the following questions:** **(2 x 10)**

a) What is a key-value store? Give one example.

b) Differentiate between NoSQL and relational databases.

c) What is the purpose of data sharding in distributed systems?

d) State the function of a reducer in the MapReduce paradigm.

e) List two features of Ceph storage.

f) What is the purpose of HDFS in cloud data management?

g) What is the function of Kubernetes in cloud deployment?

h) Define the term “virtual hard drive” in cloud storage.

i) What is archival storage, and when is it used?

j) List the characteristics of software-defined architecture.

Part-II

Q2 **Only Focused-Short Answer Type Questions- (Answer Any Eight out of Twelve)** **(6 x 8)**

a) Compare and contrast the MapReduce paradigm with modern stream processing frameworks.

b) Explain how the Paxos consensus algorithm works.

c) How do distributed hash tables (DHTs) enable efficient key-value storage?

d) Discuss the implications of Brewer’s CAP theorem in designing cloud systems.

e) Describe a cloud middleware stack and its components.

f) How does MaaS contribute to physical infrastructure automation? Explain.

g) Describe the architecture of a PaaS offering and its developer-focused advantages.

h) How do JVM-based applications containerize for cloud deployment? Explain.

i) Compare cloud object storage with traditional block storage.

j) Explain how containers and orchestration tools help manage scalable big data pipelines.

k) Compare Hive with traditional RDBMS in analyzing large-scale datasets.

l) How does a serverless architecture facilitate scalable data ingestion and transformation? Explain.

Part-III

Only Long Answer Type Questions (Answer Any Two out of Four) **(16 x 2)**

Q3	Describe a fault-tolerant distributed application using C++ that implements replication and recovery. Explain the logic and assumptions.	**(16)**
Q4	Describe a complete cloud infrastructure for a mid-sized enterprise using virtualization and container orchestration with a suitable diagram.	**(16)**
Q5	Compare the performance and management of a traditional three-tier app deployed on VMs vs. a microservices app deployed in containers.	**(16)**
Q6	Compare and contrast cloud-native vs. on-premise big data solutions in terms of architecture, cost, and elasticity.	**(16)**