Distributed computing systems are networks of connected computers that work together to solve complex tasks or process large amounts of data. In these systems, individual computers, called nodes, collaborate and communicate with each other to achieve a shared objective.

This approach allows us to deal with tasks that would be difficult or impossible to do on just one machine.

Key characteristics and components of distributed computing systems:

- Distributed Architecture: The system architecture is designed to distribute tasks and data processing across multiple nodes. Each node may have its own processing capabilities, memory, and storage.
- Network Communication: Communication between nodes is vital in distributed computing systems. Nodes need to exchange data, coordinate tasks, and share results to achieve a collective outcome.
- 3. Fault Tolerance: Distributed systems are designed to be fault-tolerant, meaning they can continue operating even if some nodes fail or experience issues. Data replication, redundancy, and backup mechanisms are often used to ensure resilience.
- Scalability: Distributed systems can scale horizontally by adding more nodes to the network. This enables them to handle increasing workloads and data volumes efficiently.
- 5. Load Balancing: Load balancing techniques ensure that tasks are evenly distributed among nodes to optimize resource utilization and improve overall performance.
- 6. Synchronization and Consistency: Maintaining data consistency across distributed nodes is a significant challenge. Techniques such as distributed locking, timestamps,

and consensus algorithms are used to synchronize data and maintain consistency.

 Distributed File Systems: Distributed file systems like Hadoop Distributed File System (HDFS) and Google File System (GFS) allow data to be stored and accessed across multiple nodes.

Examples of Distributed Computing Systems:

- Hadoop: Hadoop is an open-source distributed computing framework designed to process and analyze vast amounts of data across clusters of commodity hardware. It is widely used for big data processing, using the Hadoop Distributed File System (HDFS) and the MapReduce paradigm.
- Apache Spark: Spark is another popular distributed computing system that provides inmemory data processing capabilities, making it significantly faster than Hadoop's MapReduce for certain workloads. Spark supports various data processing tasks, including batch processing, real-time streaming, and machine learning.
- 3. Distributed Databases: Systems like Apache Cassandra and MongoDB are examples of distributed databases. They are designed to store and manage data across multiple nodes with high availability and fault tolerance.
- 4. Cloud Computing Platforms: Cloud platforms like Amazon Web Services (AWS) and Microsoft Azure offer distributed computing services, providing scalable computing resources to handle diverse workloads.
- 5. Peer-to-Peer Networks: Peer-to-peer (P2P) networks use distributed computing principles to share resources, files, or services among interconnected nodes.