

## Case Study: GreenValley Data Center's CPU Scheduling Challenge

- GreenValley Data Center provides cloud computing services and hosts a variety of applications ranging from large-scale data processing to interactive web services. They recently upgraded their server infrastructure and are now looking to optimize CPU scheduling to enhance performance and customer satisfaction.

The data center uses a mix of different CPU scheduling algorithms to handle two major types of workloads

### **Batch Processing Workload:**

1. Involves handling large volumes of data processing tasks, typically for data analytics.
2. Tasks are CPU-intensive and have longer execution times.
3. GreenValley uses a **First-Come, First-Served (FCFS)** scheduling algorithm for this.

### **Interactive Workload:**

1. Includes serving web pages, handling user queries, and other interactive applications.
2. Tasks are time-sensitive and require quick response times.
3. A **Preemptive Priority Scheduling** algorithm is used, where each task is assigned a priority based on its urgency.

**Challenges:** The major challenge GreenValley faces is balancing these two very different types of workloads efficiently. Batch processing tasks, while important, can consume considerable CPU time, potentially leading to longer response times for the interactive tasks.

## Questions for Analysis:

### First-Come, First-Served Algorithm:

1. What are the potential impacts of using the FCFS algorithm for batch processing tasks on overall system performance?
2. How might task starvation occur in this scenario, and what strategies can be employed to mitigate it?

### Preemptive Priority Scheduling:

- How does preemptive priority scheduling work, and why is it suitable for interactive workloads?

### Real-World Implications:

- What are the practical considerations in implementing these algorithms in a real-world data center environment?