

Outline

- Lottery Scheduling: Flexible Proportional Share Resource Management. C. Waldspurger, and W. Weihl - 1st OSDI
- Resource Containers: A New Facility for Resource Management in Server Systems. Gaurav Banga, Peter Druschel, and Jeffrey C. Mogul - 3rd OSDI



Problem 1

- Flexible, responsive fair share scheduling of resources
 - Changes should be reflected immediately
- Solutions: Priority based allocation
 - UNIX uses multiple priority levels with round-robin within a priority level and decay-usage
 - Lottery scheduling: Randomized resource allocation
 - Responsive
 - Proportional share
 - Modular resource management



Lottery Scheduling

- Resource rights
 - Abstract: independent of machine details
 - Relative: Relative to contention for resource
 - Uniform: Rights for multiple resource types uniformly represented as tickets
- Lotteries:
 - Probabilistically fair
 - Order clients with decreasing ticket counts
 - Tree of partial ticket sums



Modular Resource Management

- Ticket transfer: Transfer from one to another client
- Ticket inflation: Escalate rights by creating tickets
- Ticket currencies: Built on base currencies
- Ticket compensation: Fairness for IO bound tasks

- Applications:
 - Client-server
 - Multimedia
 - Servers



Problem 2: Resource allocation boundaries

- OS notion of protection domain and resource principal coincide with process abstraction
- Resource containers separate this abstraction
- Resource containers abstract resource allocation across process, kernel boundaries
- Applicable for web services etc.

