

# Threads

By Ed Terracciano

# Threads

- Thread Models
- Scheduling

# Thread Models

## 1. Kernel level threads


- a. 1 to 1
- b. Every thread is seen by the kernel.

## 2. User level threads

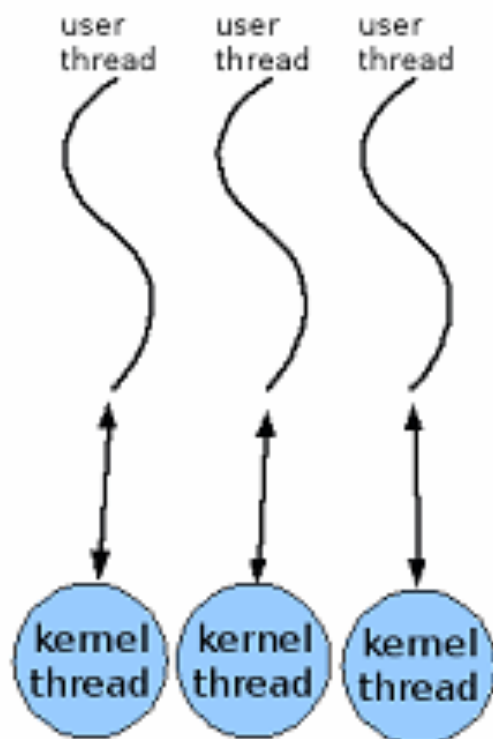
- a. n to 1
- b. Threading done within the application.

## 3. Hybrid Threads


- a. n to m
- b. Combination of both kernel and user level threads.



many to one



one to one



many to many

# Scheduling

- Co-operative Scheduling
- Pre-emptive Scheduling
  - OS balances several goals in scheduling
    1. Throughput
    2. Latency
    3. Response time
    4. Fairness
    5. Waiting time

# Scheduling Algorithms

- First in First Out
- Round Robin
- Earliest Deadline First
- Shortest Time Remaining
- Fixed Priority Preemptive Scheduling
  - Multilevel Queue Scheduling