# Supercharging PlanetLab – a High Performance, Multi-application, Overlay Network Platform

Jon Turner, Patrick Crowley, John DeHart, Amy Freestone, Brandon Heller, Fred Kuhns, Sailesh Kumar, John Lockwood, Jing Lu, Mike Wilson, Charlie Wiseman, Dave Zar
Department of Computer Science and Engineering

Supported by the National Science Foundation



# Overlay Hosting Service

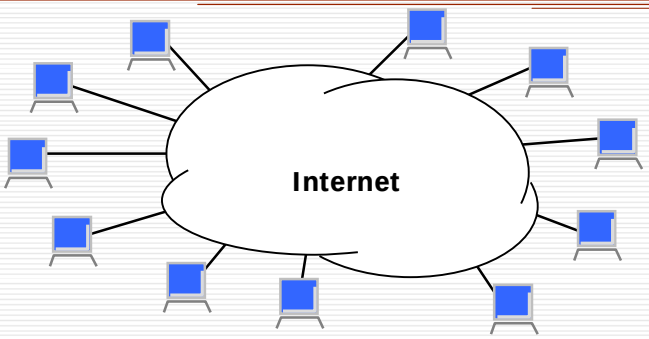


- Shared overlay infrastructure supporting many services
- Vehicle for research and deployment
- Testbed: PlanetLab

# PlanetLab

- Shared overlay network testbed
- Applications run as user-space processes in virtual machines
  - » Limited throughput
  - » High, unpredictable latency

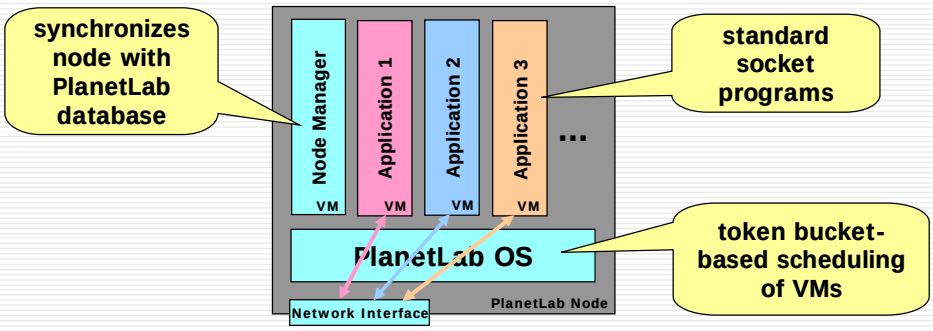





# Supercharging PlanetLab

- Leverages network processor technology
- Standard fast-path/slow-path application structure
- Removes performance limitations
  - » Supports Internet-scale throughput
  - » Supports latency-sensitive applications
- Allows existing PlanetLab applications to run unmodified

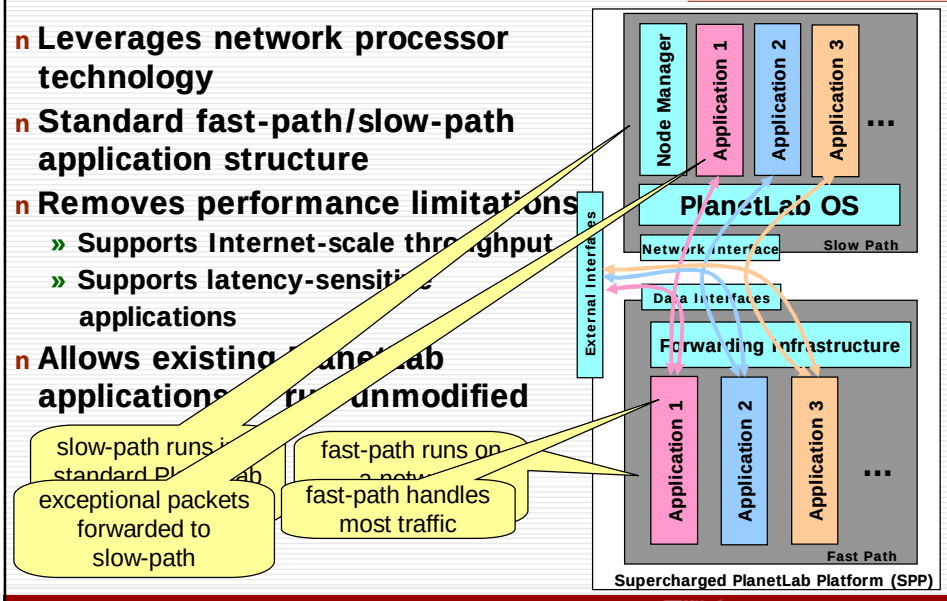

# SPP Components

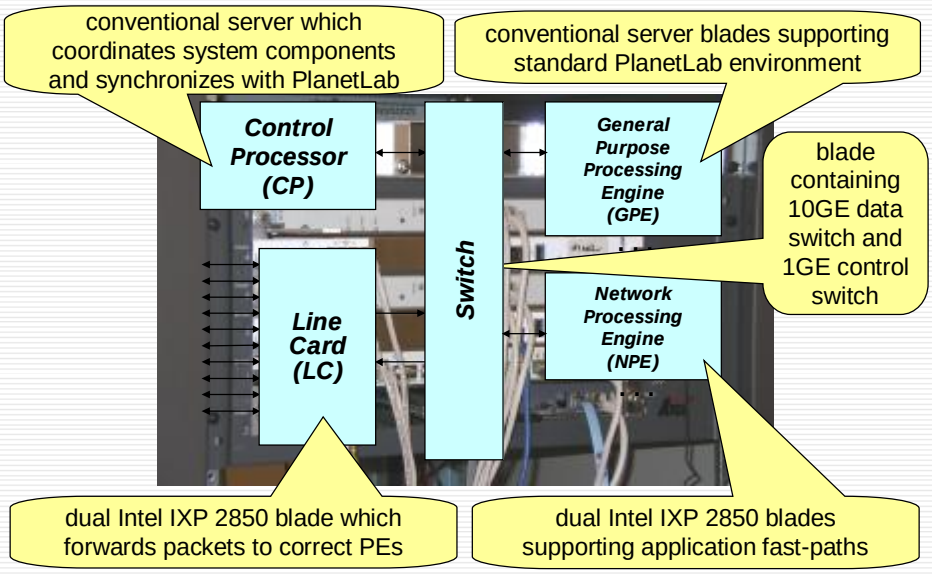




# IXP 2850 Overview

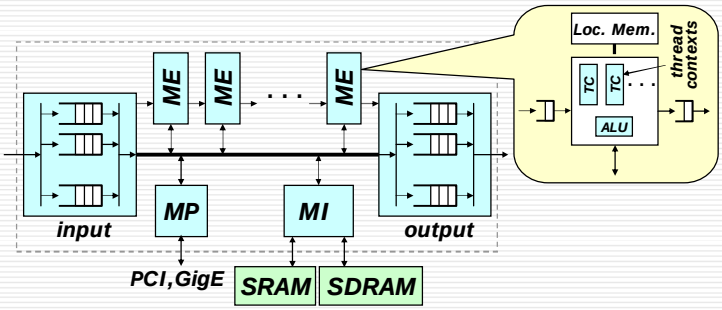


- 16 multi-threaded MicroEngines (MEs)
  - 8 thread contexts with rapid switching capability
  - Fast nearest-neighbor connections for pipelined apps
- 3 SDRAM and 4 SRAM channels (optional TCAM)
- Management Processor (MP) for control

# System Control

- Instantiate new application
- Open socket
- Instantiate fast-path

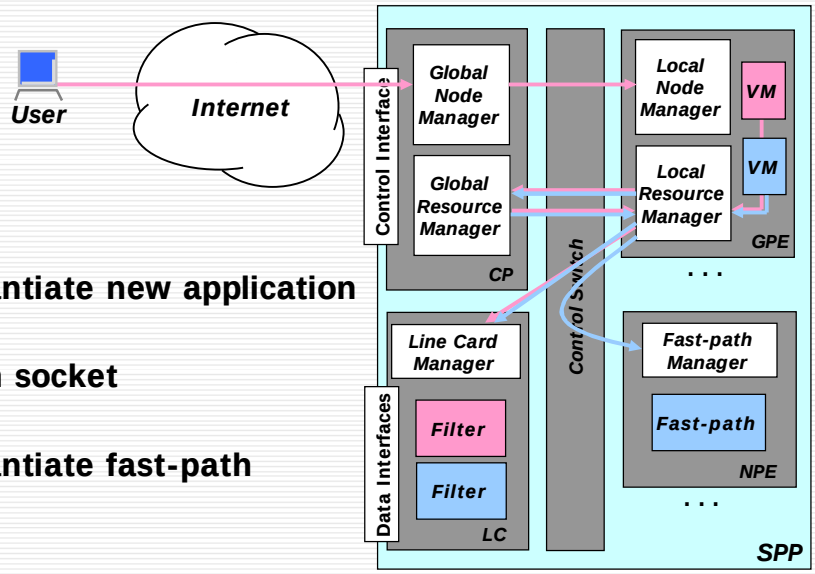




# Sharing the NPE

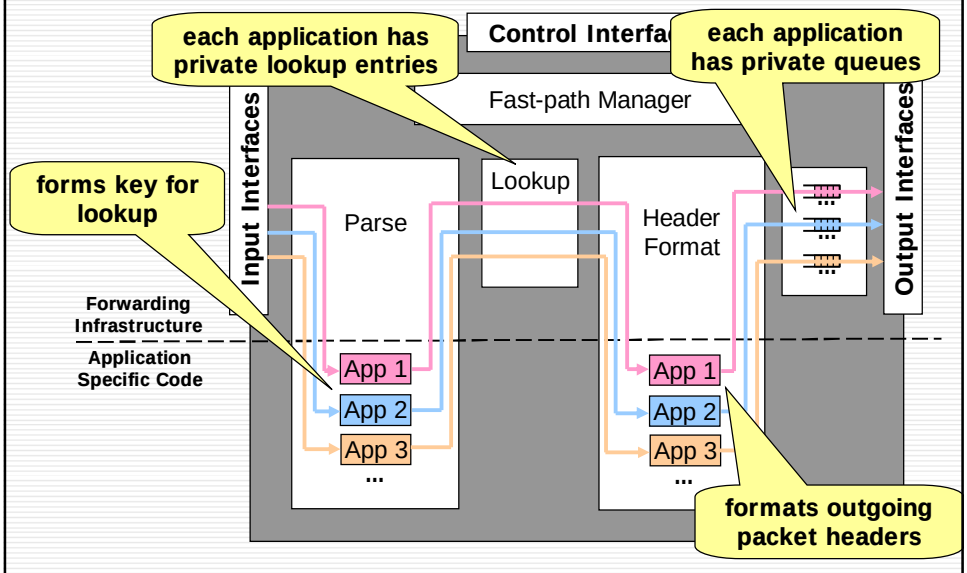

# Evaluation

- IPv4
  - Packets arrive/depart in UDP tunnels
- Internet Indirection Infrastructure (i3)
  - Packets contain *triggers* matched to IP addresses
  - No match at local node results in Chord forwarding

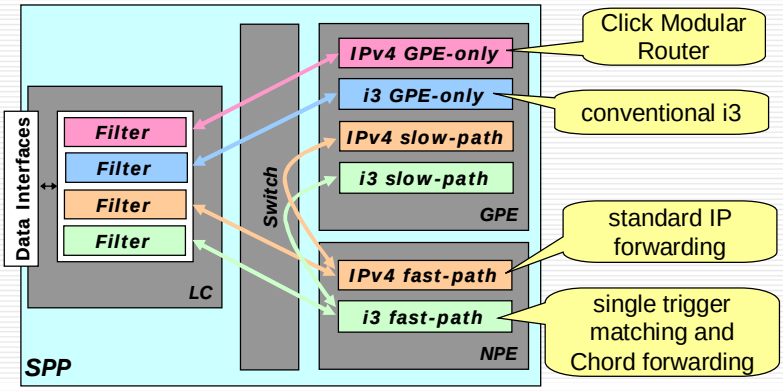





# IPv4 Throughput Comparison

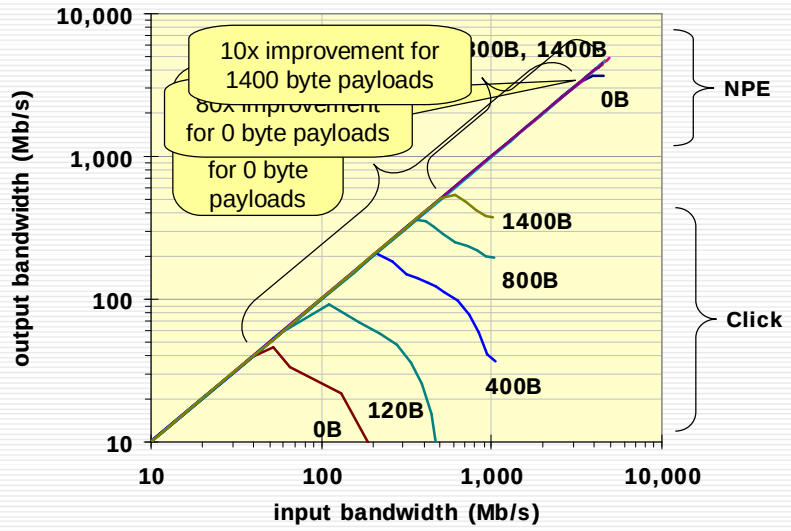

# IPv4/i3 Fast-Path Throughput Comparison

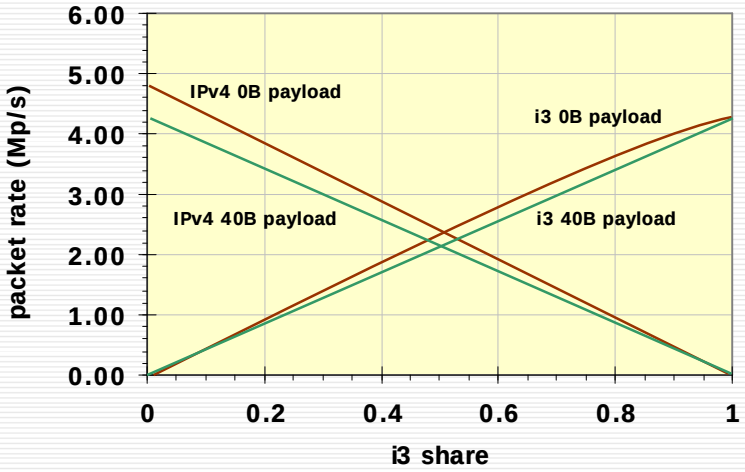


- Constant input rate of 5 Gb/s



# IPv4 Latency Comparison

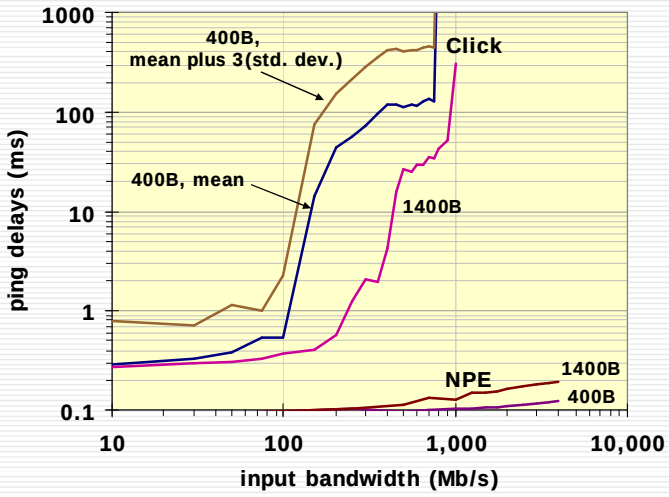


- 8 IPv4 instances

# Summary

- Base platform intended for overlay hosting service
- An SPP node removes performance limitations found in conventional PlanetLab nodes
- Standard fast-path/slow-path application structure eases deployment
- Future work includes
  - More flexible IXP-based NPE implementation
  - NPEs built on other hardware
  - Automatic NPE code verification
- Targeting 2 SPP nodes available by end of 2007

- Questions?



# i3 Latency Comparison

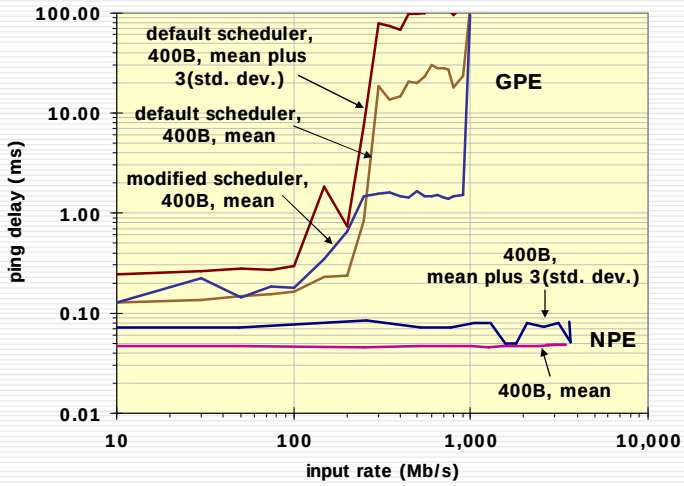


- Two i3 instances for background traffic
- Measurement traffic traverses separate instance