

Welcome

To Advance through Presentation  
Use Page Up and Page Down Keys



99 | Worldwide  
Developers  
Conference



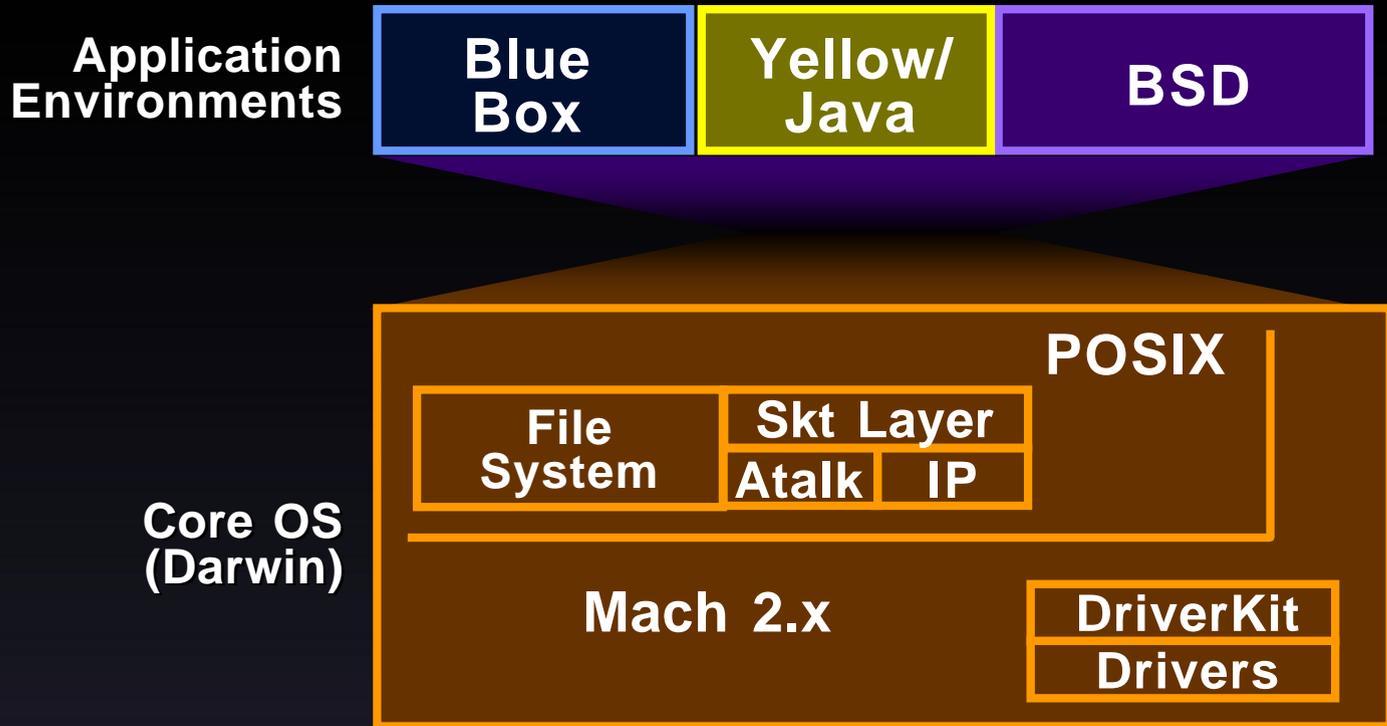
99 | Worldwide  
Developers  
Conference

# Mac OS X Kernel

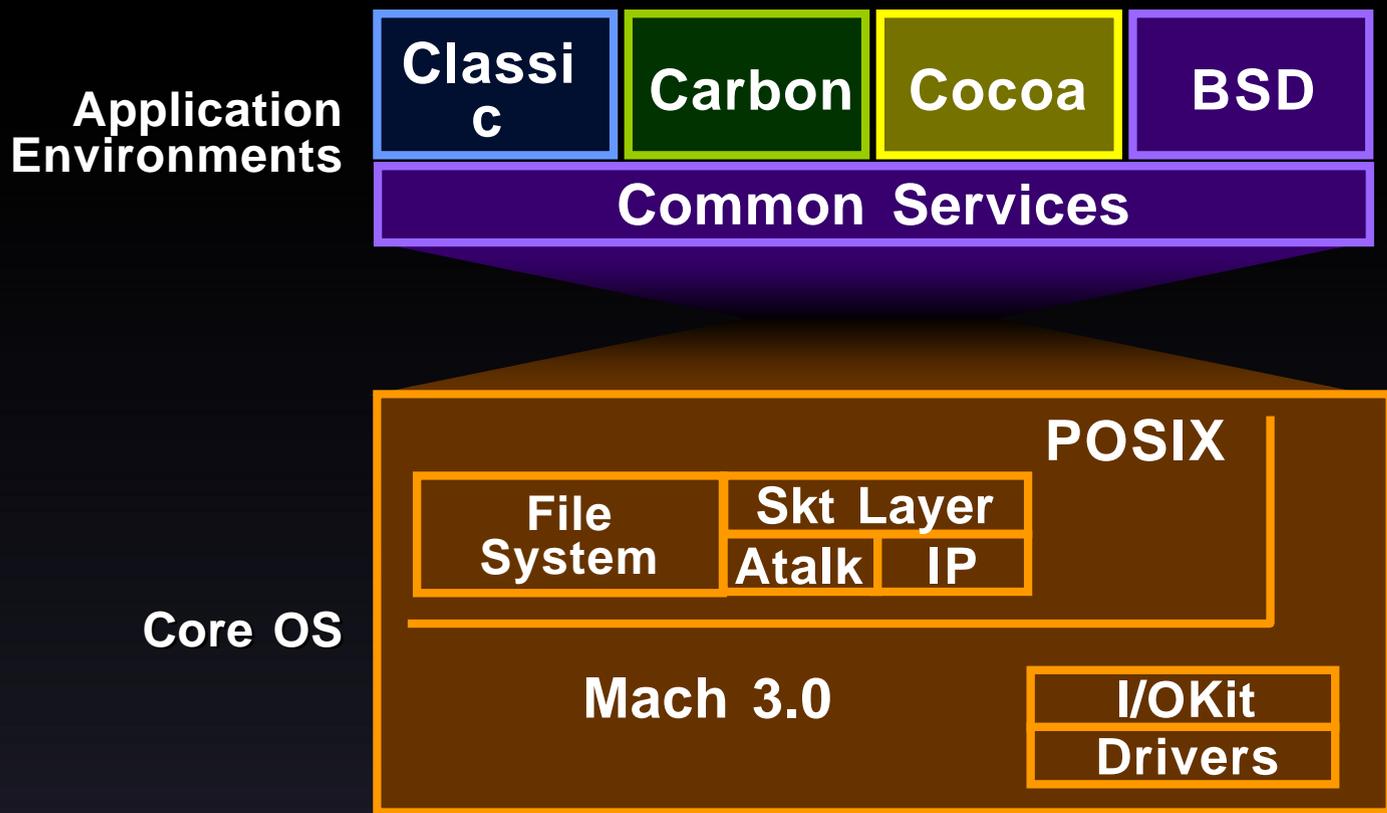
Brett Halle

Manager,  
Core OS Engineering

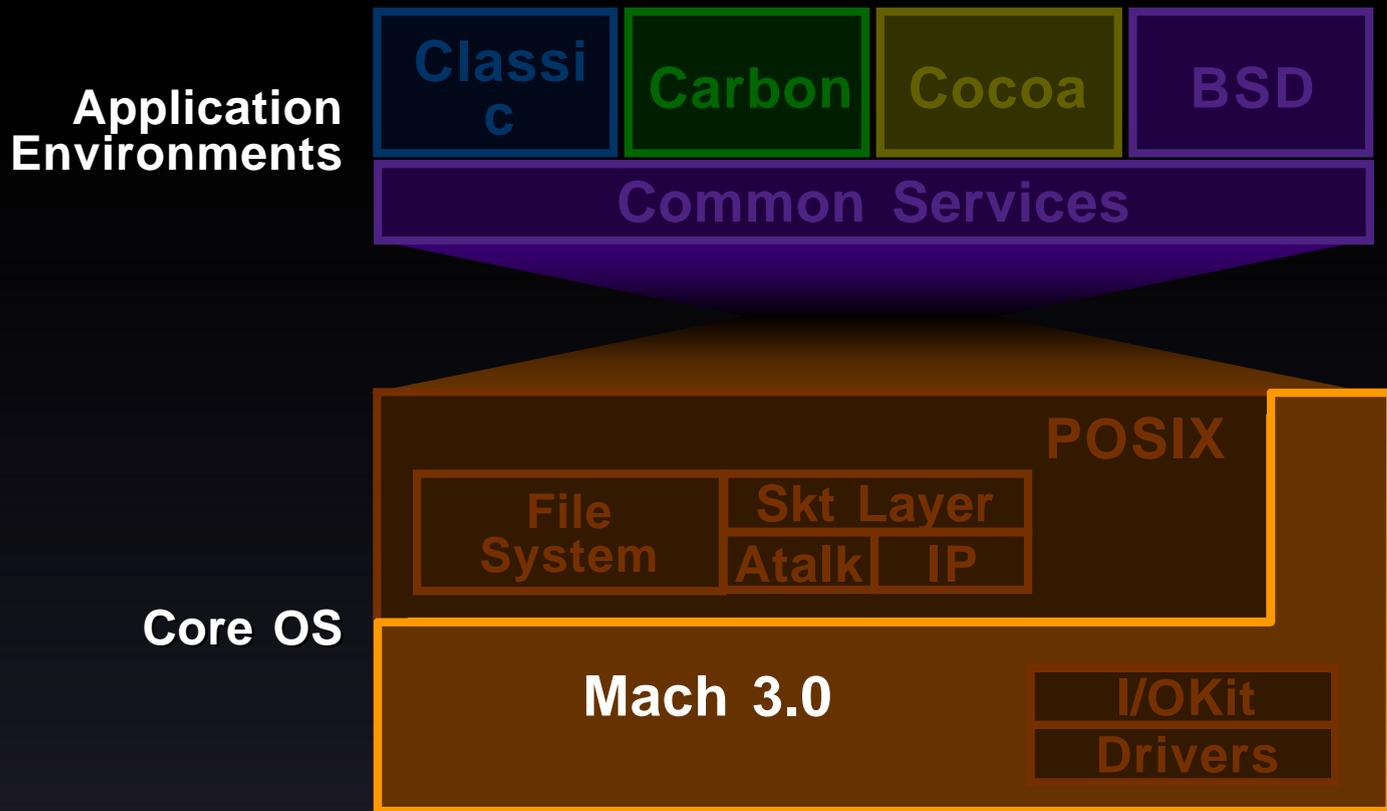
# Mac OS X Server



# Core OS and Mac OS X



# Core OS—Mach



# Mach

- The “foundation” of the OS
- Responsible for:
  - Preemption / scheduling
  - Memory protection / management
  - Low level task communication
- NOT responsible for:
  - I/O, networking, file systems, ownership, application API's



# Mach

- A little history . . .
- Mac OS X Server: based on CMU's Mach 2.x plus a handful of Mach 3.0 features
- Mac OS X: based on Mach 3.0
  - Originally developed at CMU
  - Significant work done at The Open Group Research Institute
  - Further work done at Apple (MkLinux)



# Mach 2.x

- Typed IPC
- Monolithic architecture
  - VM system (pager)
  - BSD, I/O, etc.
- Limited or no SMP capability
- Real-time not a consideration



# Mach 3.0

- Our version derived from The Open Group Research Institute's "MK 7.3"
- Better performance
  - Kernel preemption
  - Significantly improved VM system
  - IPC and RPC enhancements
    - For example, untyped IPC
  - Better real-time support



# Mach 3.0

- More modular architecture
  - Based on components
- Scheduling framework
- Processor neutral
- MP enabled
- In-line kernel debugger



# Mach Philosophy

- Simple, extensible communication kernel
- Object oriented
  - Object reference:  
communication channel
    - Also unit of protection
- Synchronous and asynchronous IPC
  - Client / Server
  - Exceptions



# Kernel Communication

- Port: communication channel
  - Accessible only via send/receive capability
  - Finest granularity of protection
- Message: collection of data objects
  - Mach 2.x: fully typed
  - Mach 3.0: untyped data, typed port capabilities



# Kernel Resources

- Task: resource container
  - Address space
  - Port rights
- Thread: lightweight unit of execution
- Memory Object: unit of backing storage
- Hardware and management abstractions
  - Host, devices, processors, components
  - Default memory manager



# Kernel Functions

- Task and Thread Management
- Virtual Memory Management
- Inter-task Communication
- Resource Management



# Tasks

- Process = Task plus one or more Threads
- Task contains:
  - Address space . . . may be sparse
    - Portions shared via inheritance or other means
  - Collection of system resources
    - References by ports
  - Some number of threads

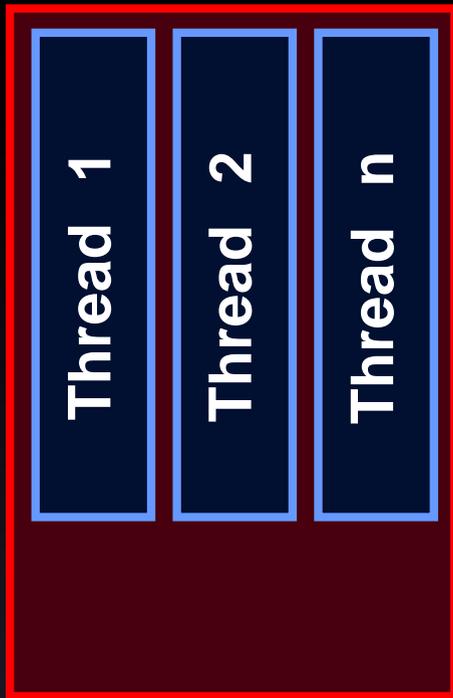


# Threads

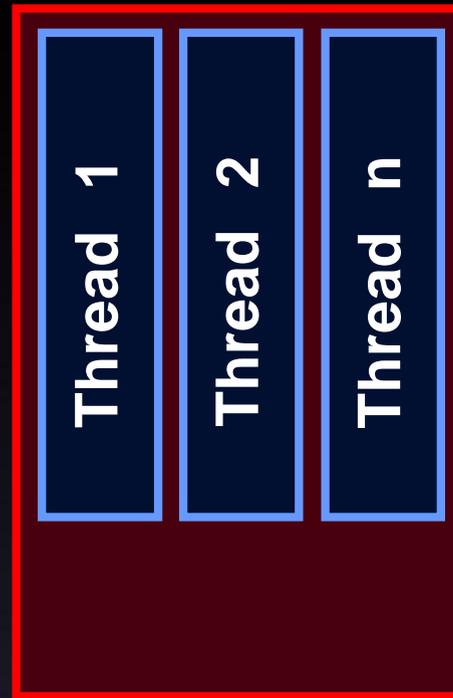
- Flow of control in task
- Has access to entire task
- Concurrent execution
  - Parallel on MP hardware
- Limited state
  - Registers plus a few other things



# Tasks and Threads



Task 1



Task 2



# Virtual Memory

- Protected address space for each task
- Flexibility in address space usage
- Controlled sharing between address spaces
- Paging implementation outside kernel (3.0)
- Copy-on-write and lazy evaluation
  - IPC integration

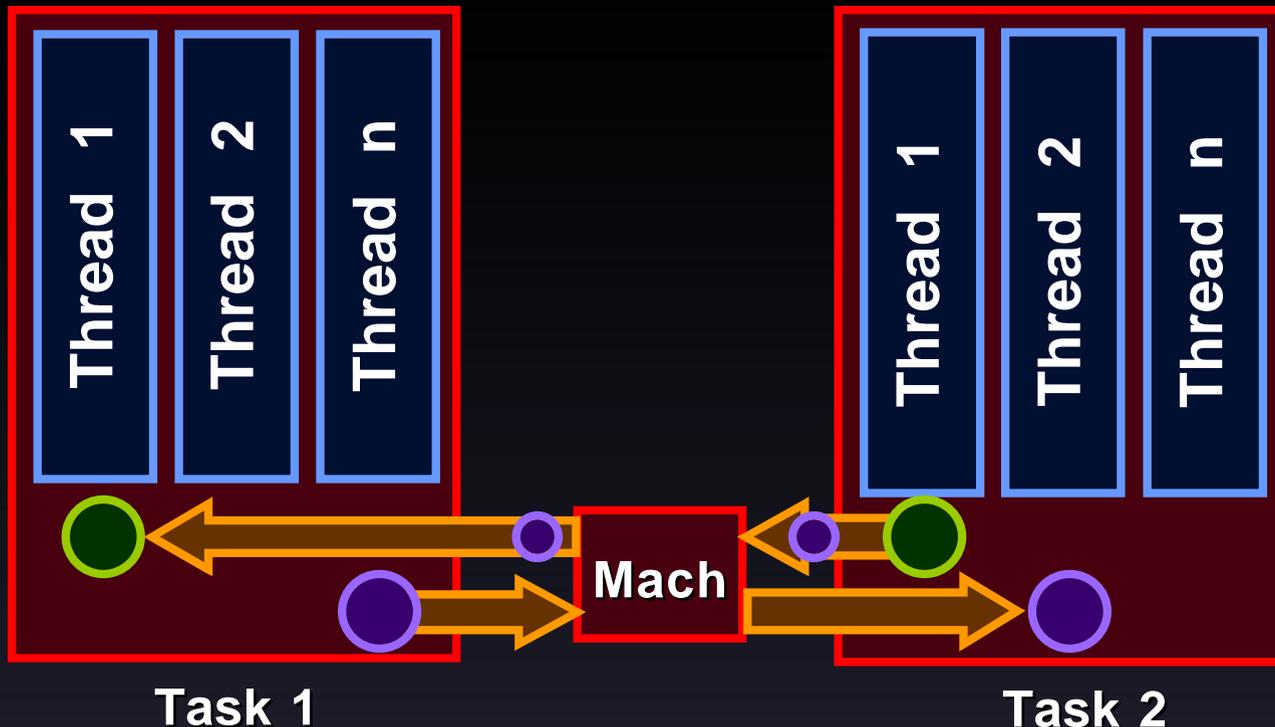


# Inter-Task Communication

- Communication via ports
  - Kernel protected objects
- Port capability = Right
  - Can only be transferred in messages
  - Port rights are **not** global
- RPC implementations
  - Optimized IPC for RPC case
  - Native RPC



# Inter-Process Communication



# Port Rights

- Single receive right
  - All threads in task may receive
  - Receive right can be moved
- Port sets
- Multiple send rights
- Send-once optimizes IPC for RPC case
- Mach 3.0 adds more port flavors



# Port Usage

- Control of kernel resources
  - Tasks
  - Threads
  - Exceptions
    - Interposing (e.g., debugger support)
- Typed vs. untyped messages (2.x vs. 3.0)
- MIG
  - Higher level abstractions recommended

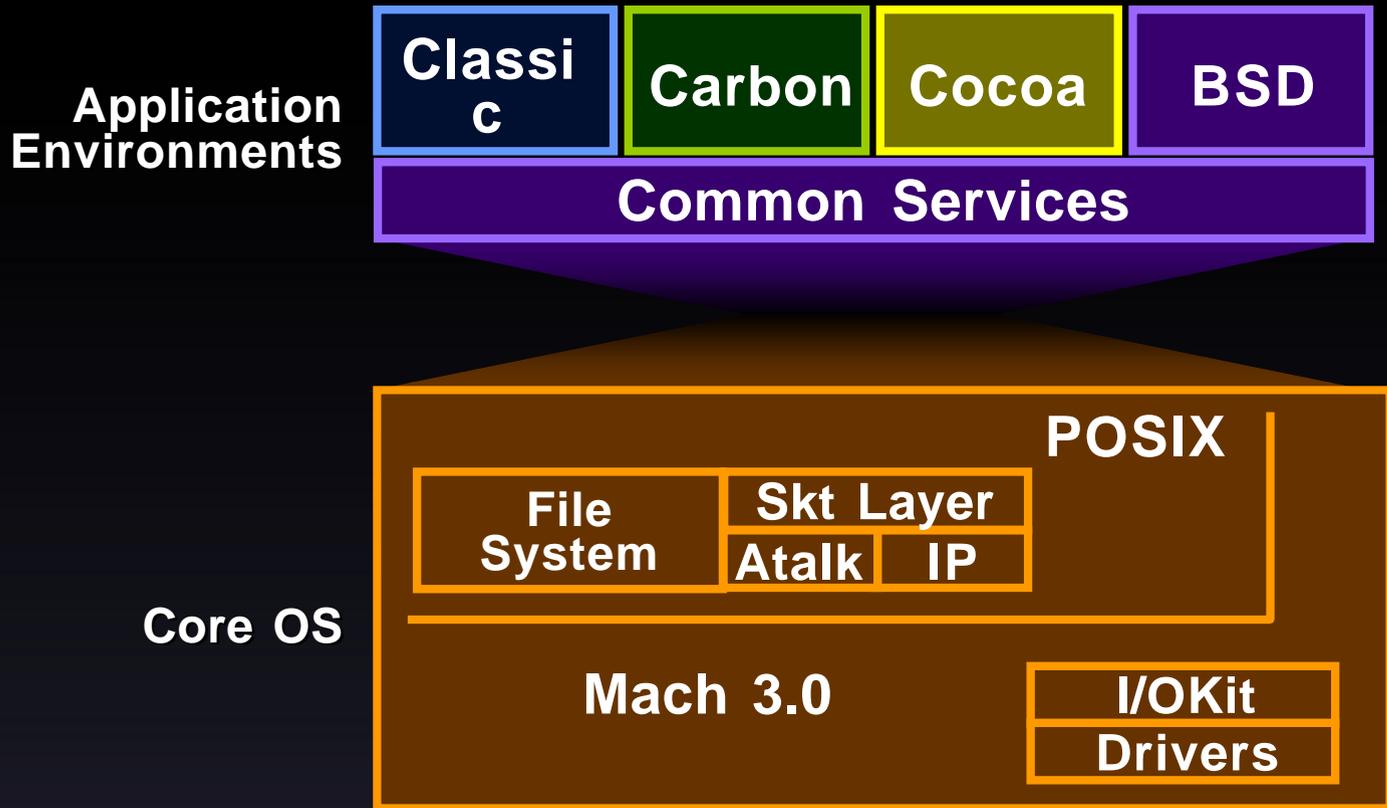


# Resource Management

- Memory
  - Address space
  - Backing store
    - Adaptive backing store
- Processor(s)
  - Access
  - Scheduling
- Ports



# Core OS and Mac OS X



# Where To Get More Info . . .

- *Programming Under Mach*  
by Joseph Boykin, et al.
- Mach 2.x source: Darwin
- Mach 3.0 source
  - MkLinux: “[www.mklinux.apple.com](http://www.mklinux.apple.com)”
  - Darwin later this year





99 | Worldwide  
Developers  
Conference

Q&A



Think different.<sup>TM</sup>



Welcome

To Advance through Presentation  
Use Page Up and Page Down Keys



99 | Worldwide  
Developers  
Conference