

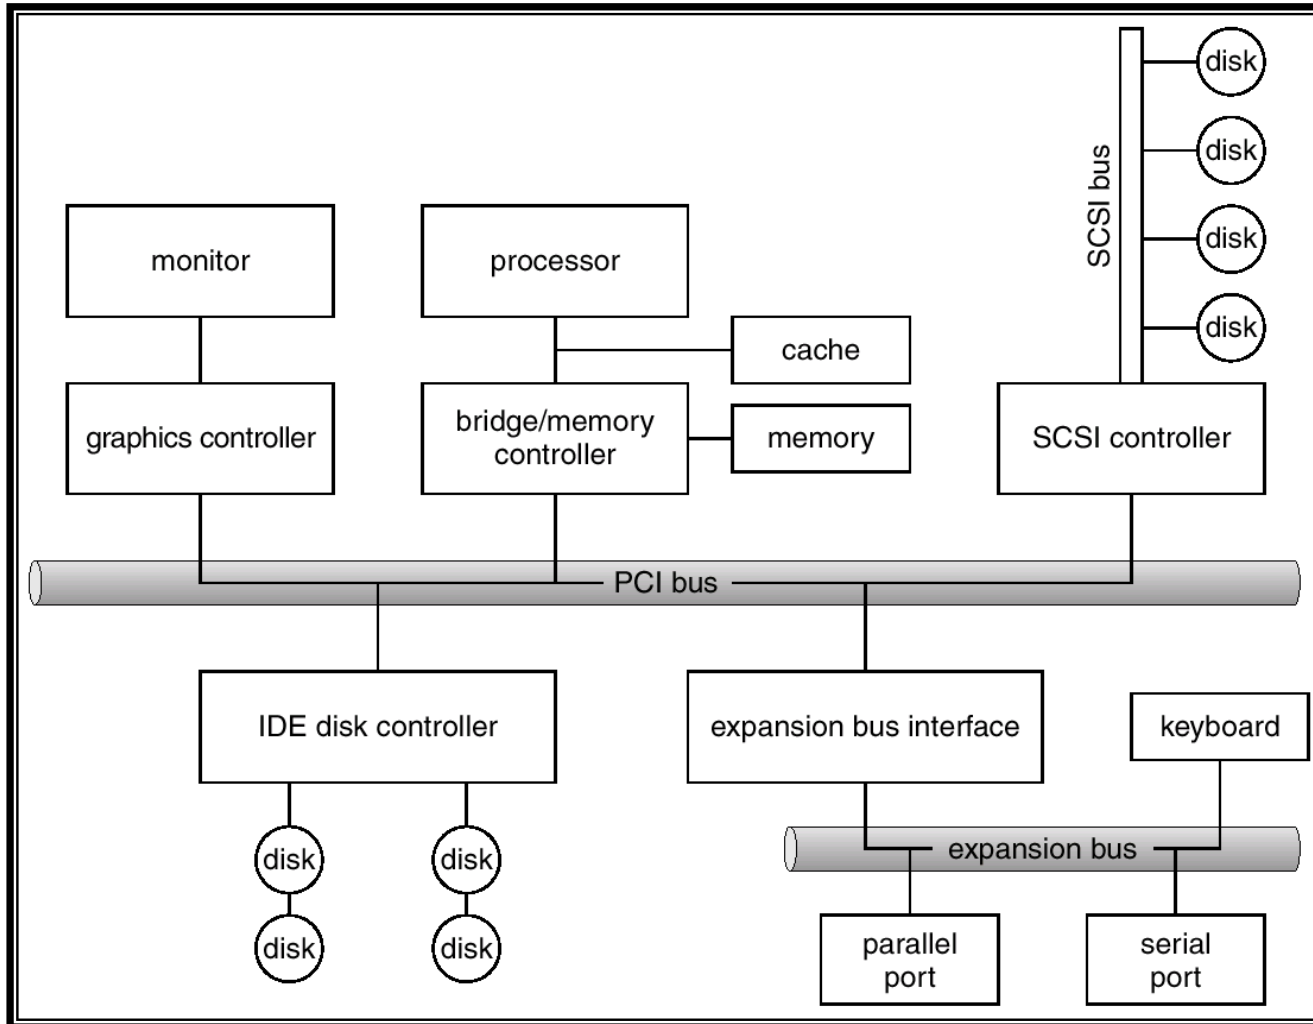
Application I/O interface – Kernel I/O subsystem

- I/O Hardware
- Application I/O Interface
- Kernel I/O Subsystem
- Transforming I/O Requests to Hardware Operations

I/O Hardware

- Categories of I/O devices
 - Human Readable-Suitable for communicating with Computer user Eg:Printer,Video Display terminals, Keyboard
 - Machine Readable-Its for Electronic Equipment
 - Eg: Disks,Tape,Controllers
 - Communication-Its with Remote Device.
 - Eg; Digital line drivers and modems
- I/O instructions control devices
- Devices have addresses, used by
 - Direct I/O instructions
 - Memory-mapped I/O

A Typical PC Bus Structure

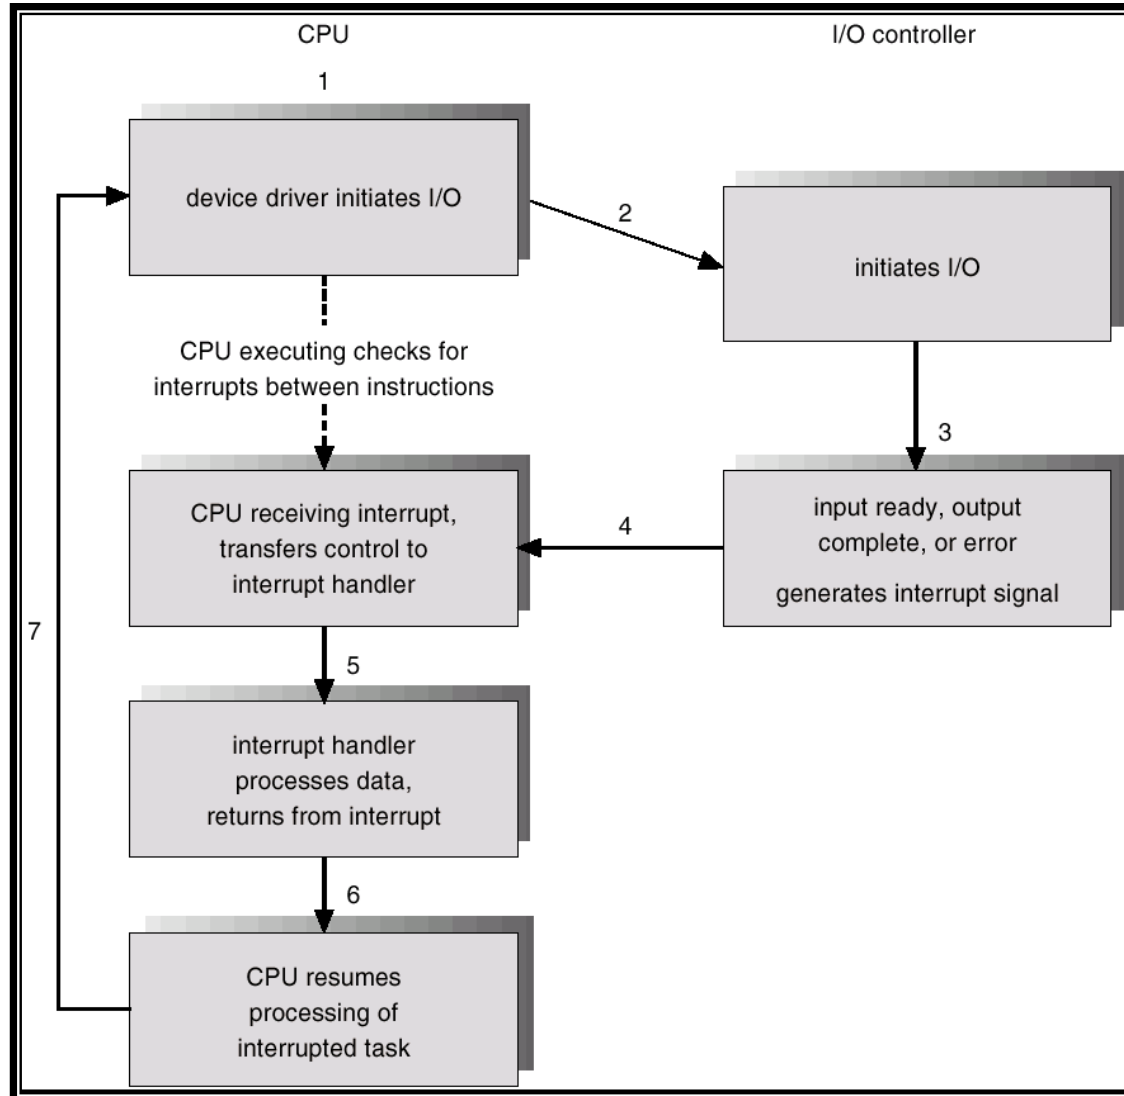


Application I/O interface – Kernel I/O subsystem

Interrupts

- CPU Interrupt request line triggered by I/O device
- Interrupt handler receives interrupts
- Maskable to ignore or delay some interrupts
- Interrupt vector to dispatch interrupt to correct handler
 - Based on priority
 - Some unmaskable
- Interrupt mechanism also used for exceptions

Interrupt-Driven I/O Cycle

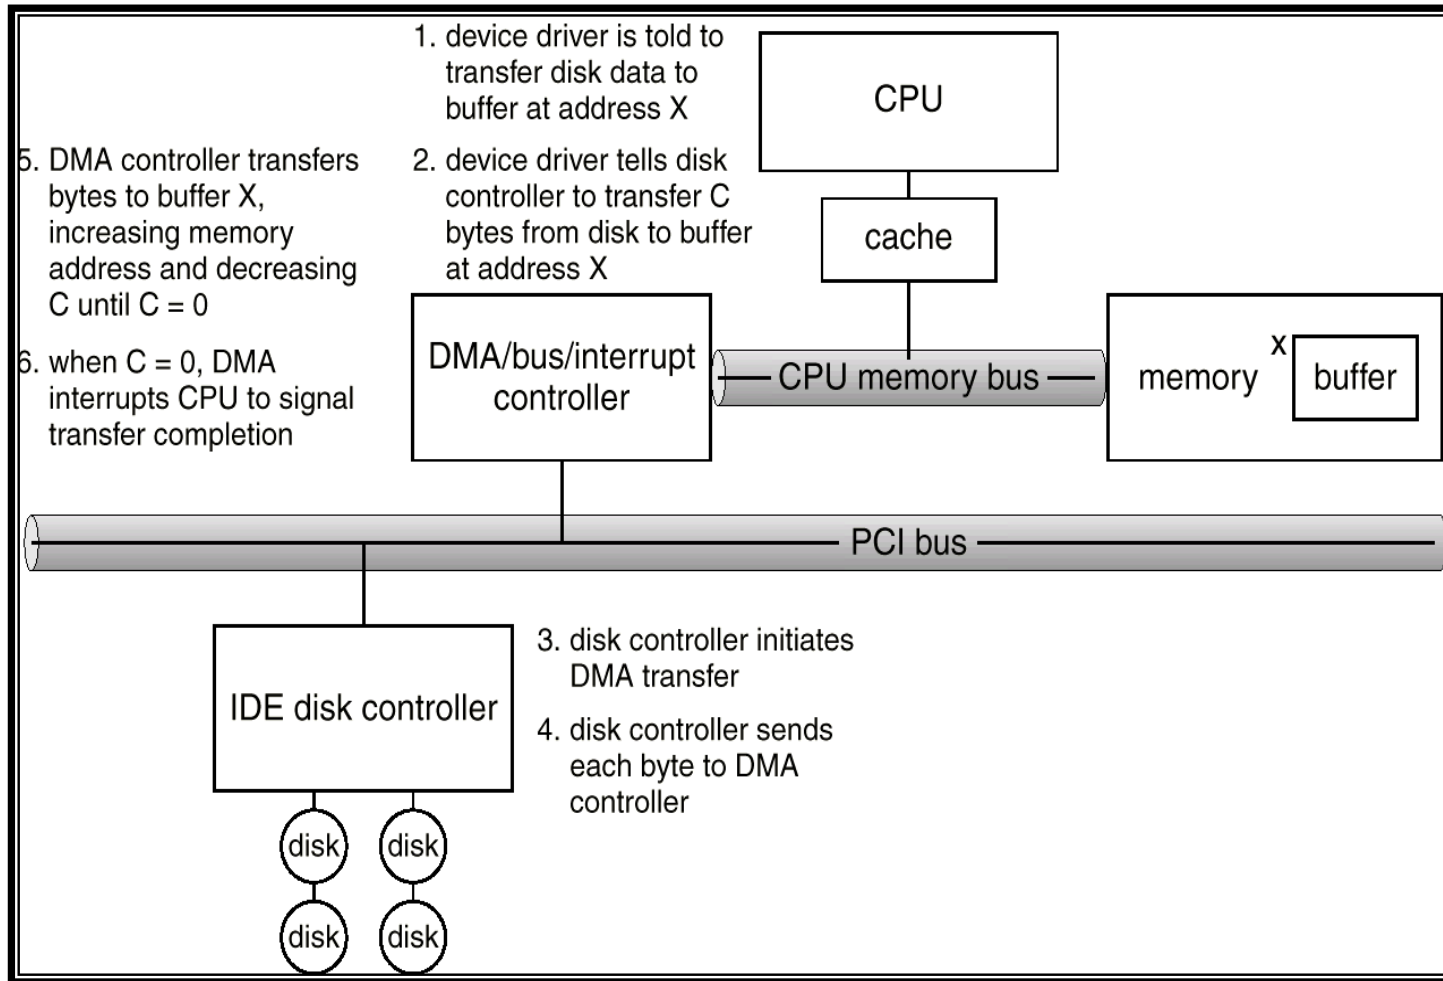


Application I/O interface – Kernel I/O subsystem

Direct Memory Access

- Special control unit provided to allow transfer of a block of data directly between an external device and the main memory, without continuous intervention by the processor.
- Used to avoid programmed I/O for large data movement
- Requires DMA controller
- Bypasses CPU to transfer data directly between I/O device and memory

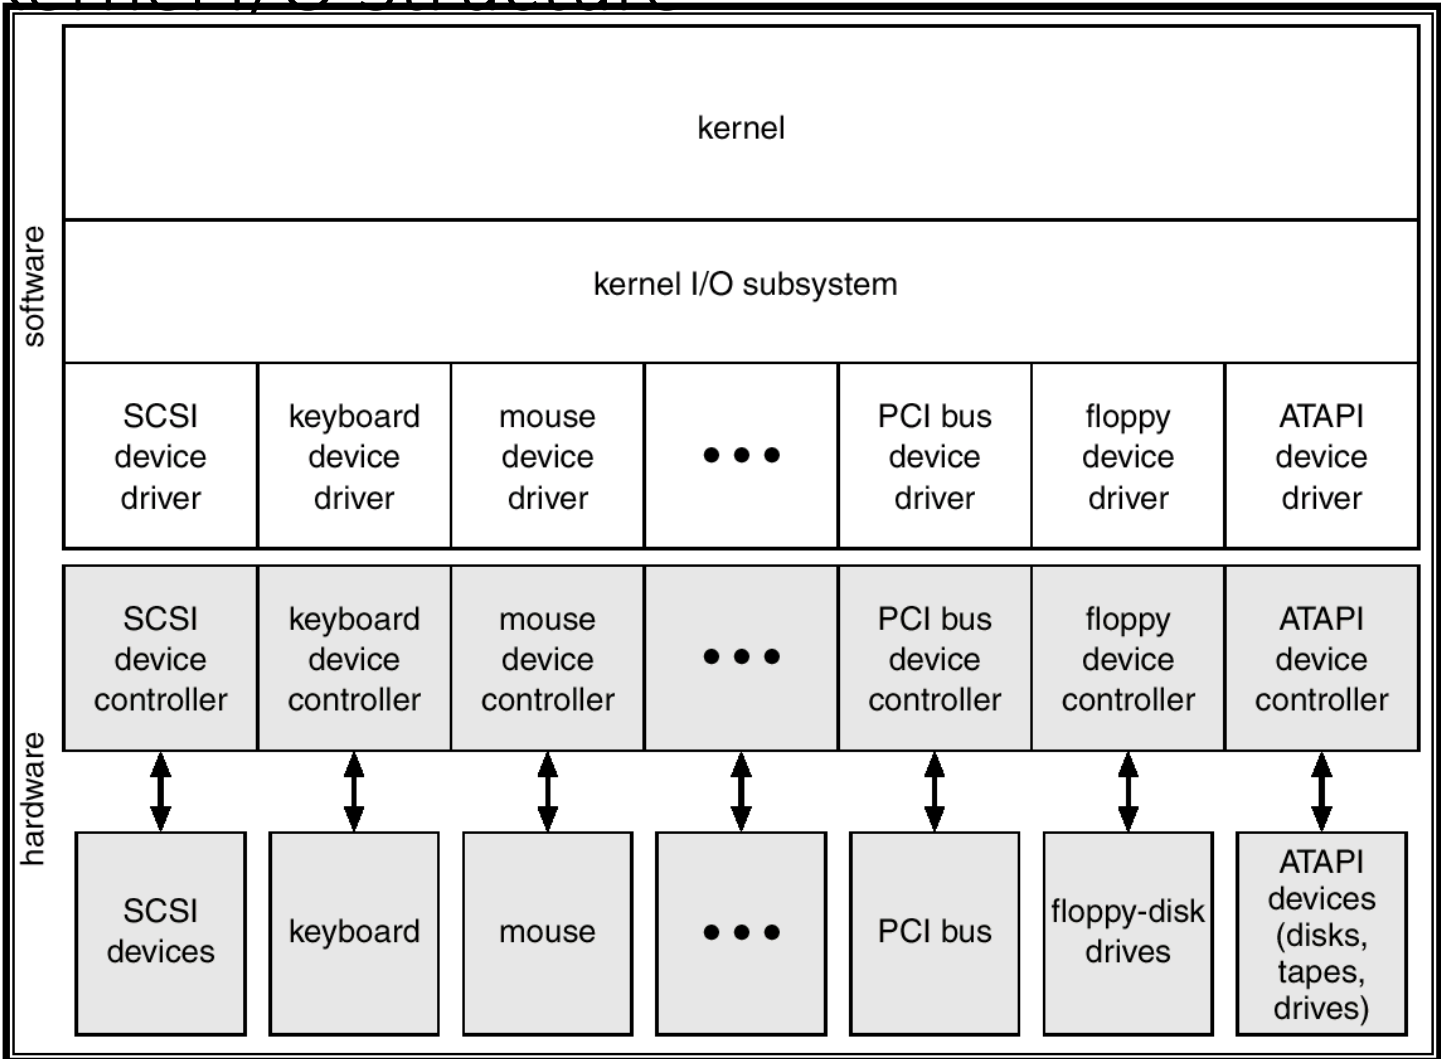
Six Step Process to Perform DMA Transfer



Application I/O Interface

- I/O system calls encapsulate device behaviors in generic classes
- Device-driver layer hides differences among I/O controllers from kernel
- Devices vary in many dimensions
 - Character-stream or block
 - Sequential or random-access
 - Sharable or dedicated
 - Speed of operation
 - read-write, read only, or write only

A Kernel I/O Structure



Block and Character Devices

- Block devices include disk drives
 - Commands include read, write, seek
 - Raw I/O or file-system access
 - Memory-mapped file access possible
- Character devices include keyboards, mice, serial ports
 - Commands include `get`, `put`
 - Libraries layered on top allow line editing

Clocks and Timers

- Provide current time, elapsed time, timer
- If programmable interval time used for timings, periodic interrupts
- `ioctl` (on UNIX) covers odd aspects of I/O such as clocks and timers

Blocking and Nonblocking I/O

- Blocking - process suspended until I/O completed
 - Easy to use and understand
 - Insufficient for some needs
- Nonblocking - I/O call returns as much as available
 - User interface, data copy (buffered I/O)
 - Implemented via multi-threading
 - Returns quickly with count of bytes read or written
- Asynchronous - process runs while I/O executes
 - Difficult to use
 - I/O subsystem signals process when I/O completed

Kernel I/O Subsystem

- Scheduling
 - Some I/O request ordering via per-device queue
 - Some OSs try fairness
- Buffering - store data in memory while transferring between devices
 - To cope with device speed mismatch
 - To cope with device transfer size mismatch
 - To maintain “copy semantics”

Kernel I/O Subsystem

- Caching - fast memory holding copy of data
 - Always just a copy
 - Key to performance
- Spooling - hold output for a device
 - If device can serve only one request at a time
 - i.e., Printing
- Device reservation - provides exclusive access to a device
 - System calls for allocation and deallocation
 - Watch out for deadlock

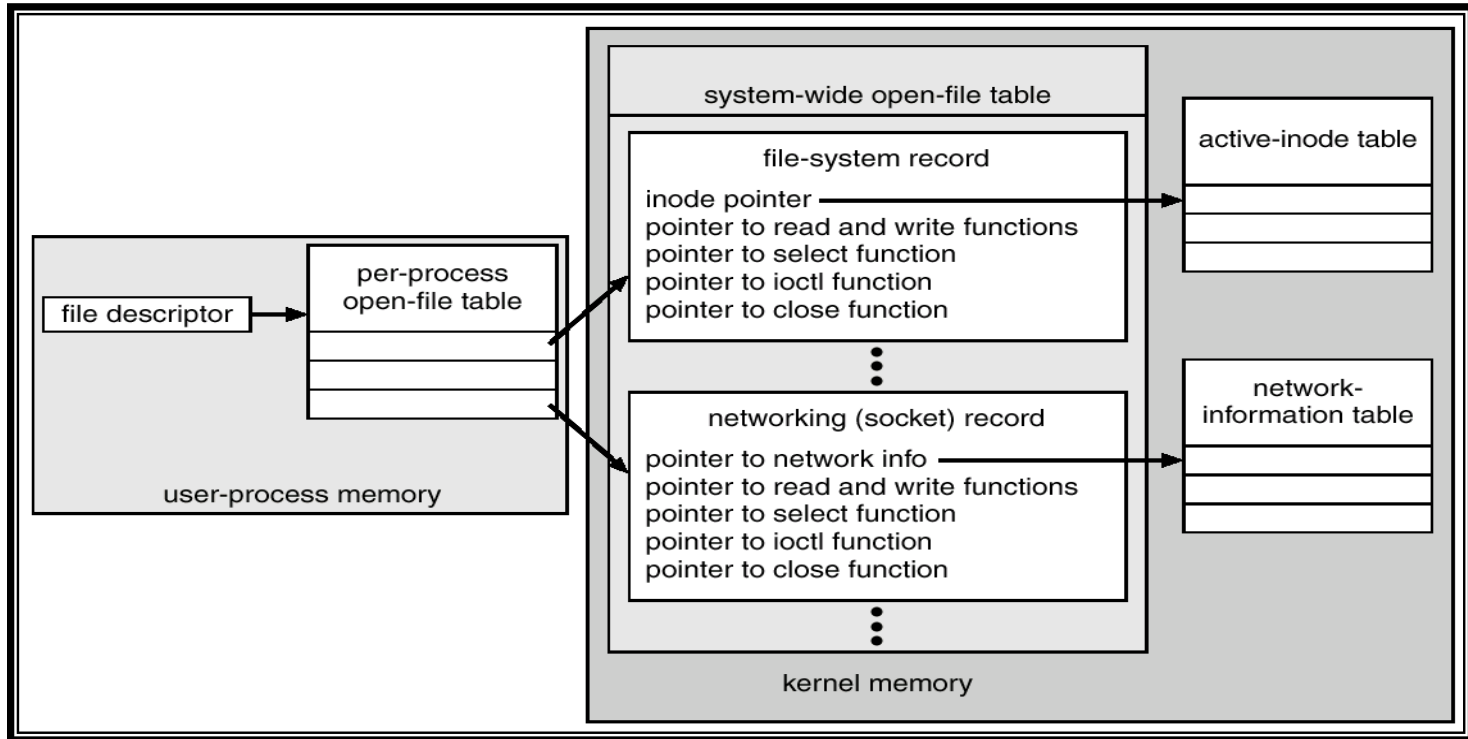
Error Handling

- OS can recover from disk read, device unavailable, transient write failures
- Most return an error number or code when I/O request fails
- System error logs hold problem reports

Kernel Data Structures

- Kernel keeps state info for I/O components, including open file tables, network connections, character device state
- Many, many complex data structures to track buffers, memory allocation, “dirty” blocks
- Some use object-oriented methods and message passing to implement I/O

UNIX I/O Kernel Structure



I/O Requests to Hardware Operations

- Consider reading a file from disk for a process:
 - Determine device holding file
 - Translate name to device representation
 - Physically read data from disk into buffer
 - Make data available to requesting process
 - Return control to process

Life Cycle of An I/O Request

