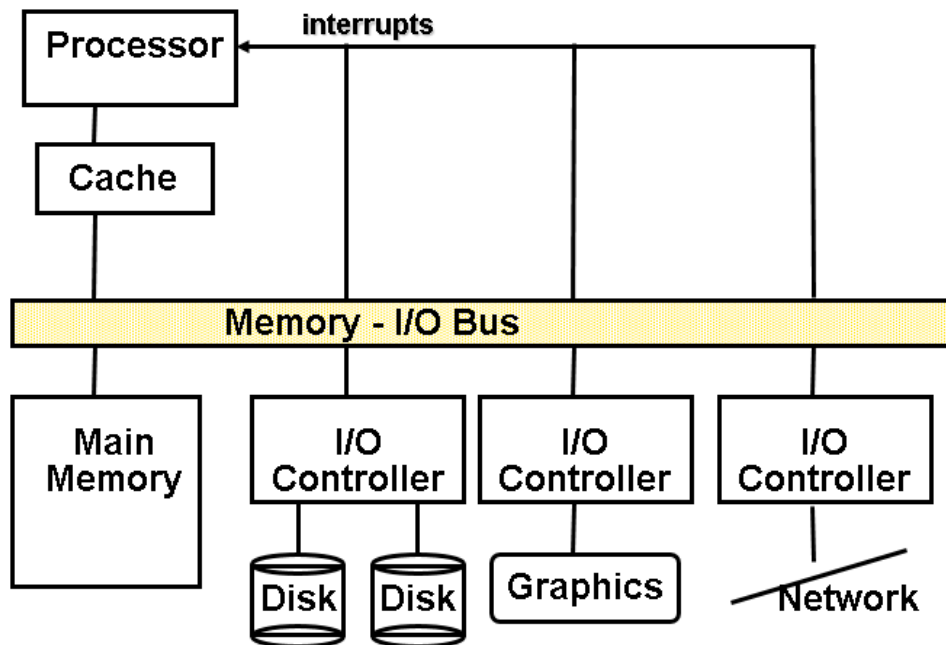
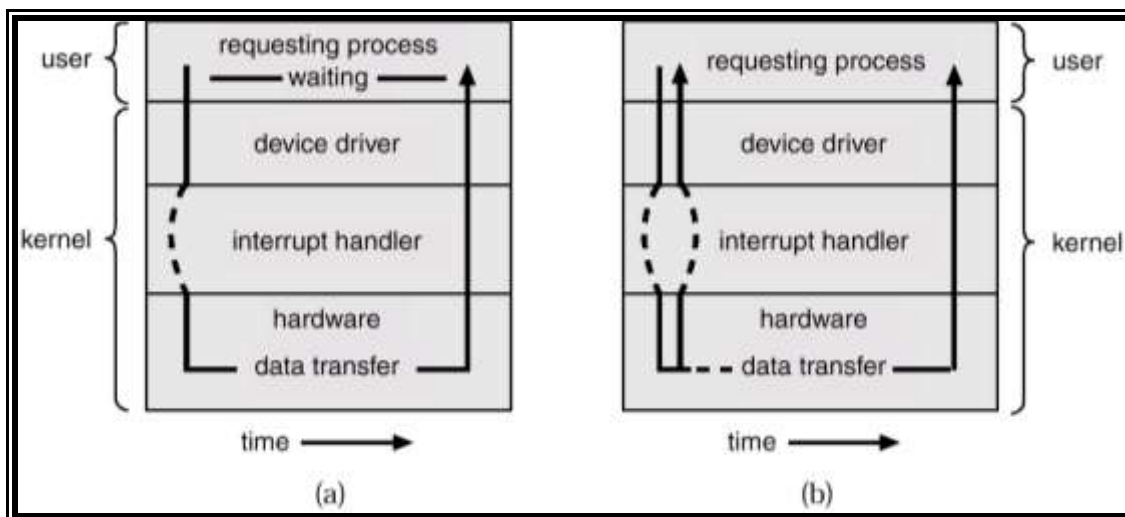


I/O Systems



I/O Device controller informs CPU that it has finished its operation by causing an **interrupt**.

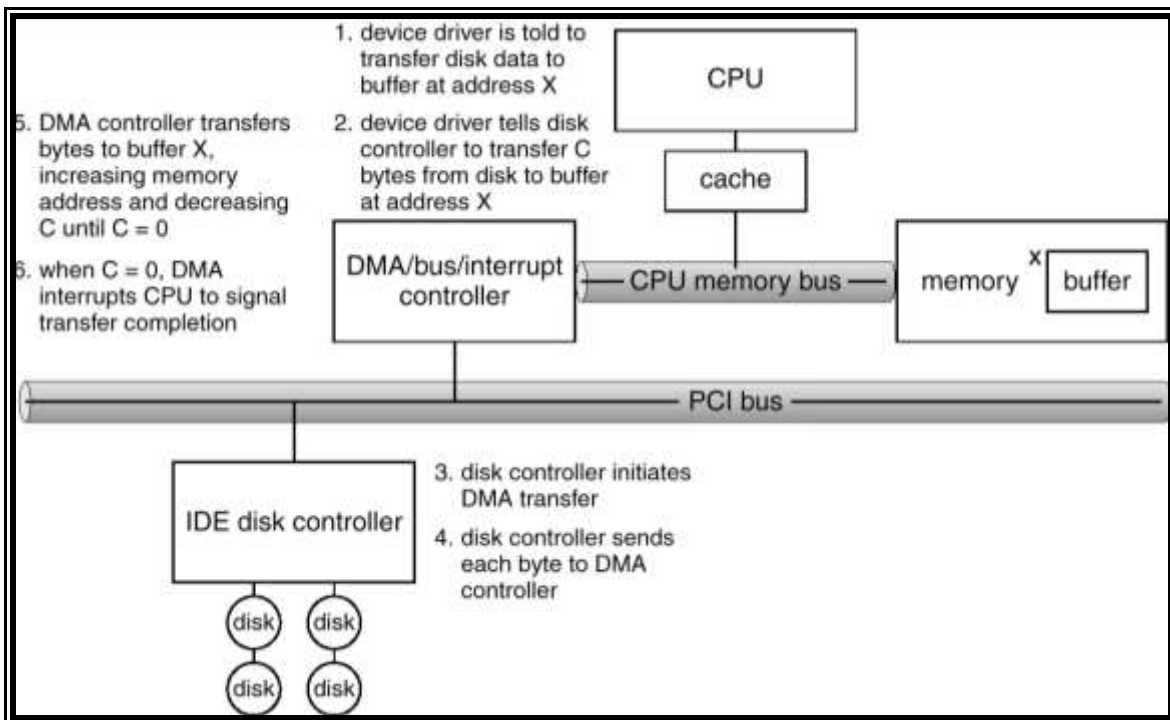
1. Two I/O Methods



Synchronous

Asynchronous

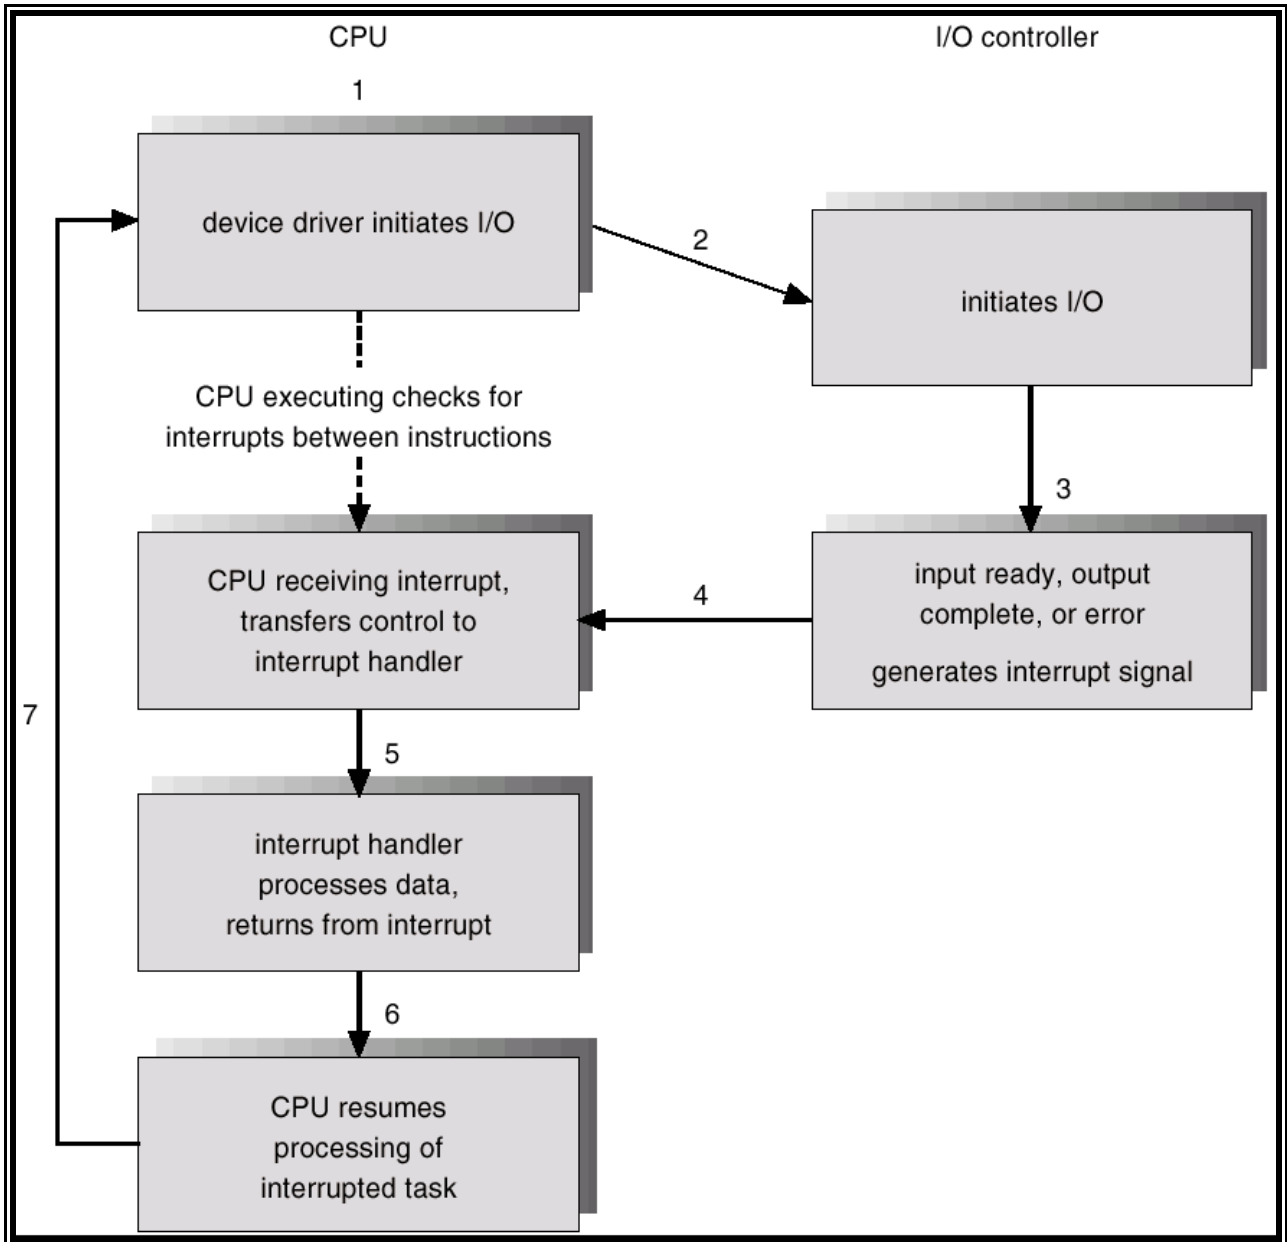
2. Direct Memory Access (DMA) : The device controller transfers blocks of data from buffer storage directly to main memory without CPU intervention.



Characteristics of I/O Devices

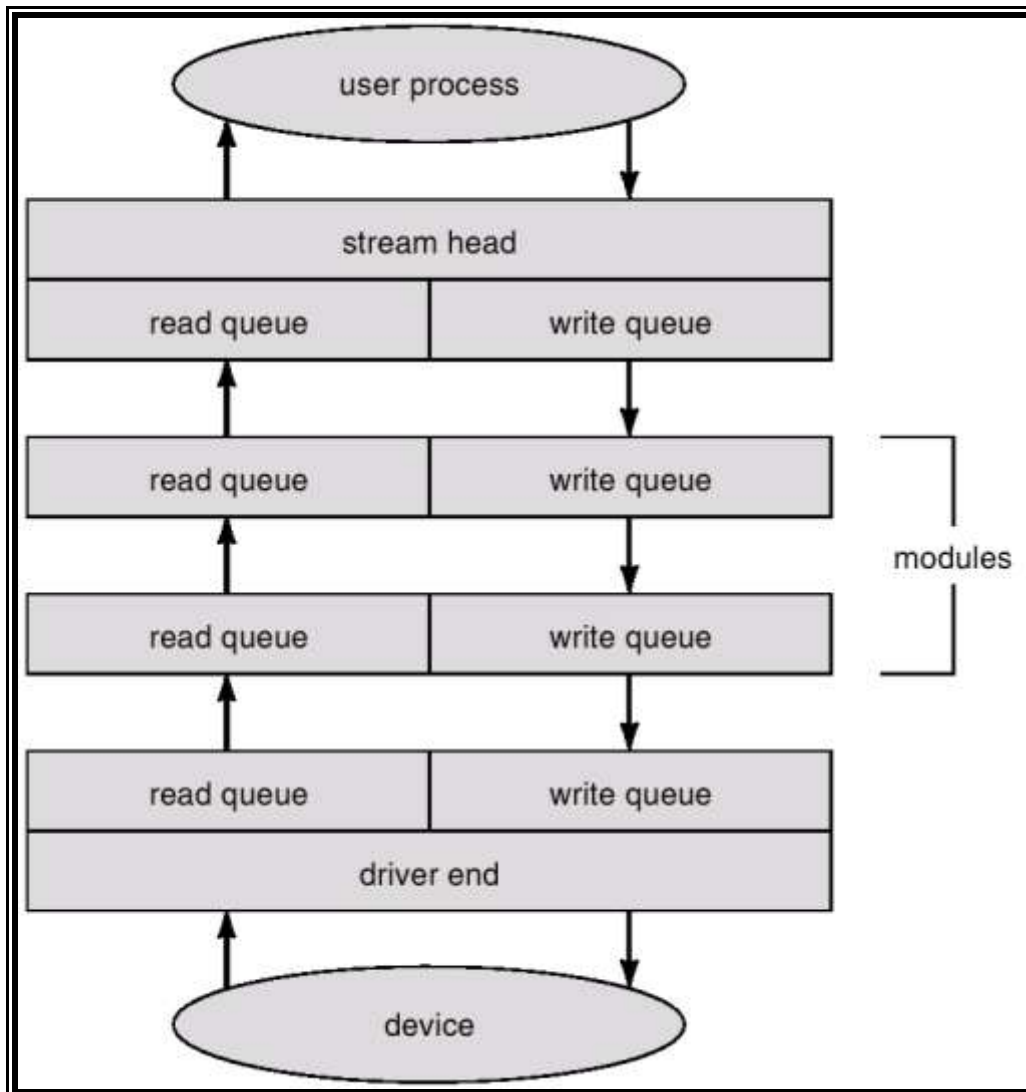
aspect	variation	example
data-transfer mode	character block	terminal disk
access method	sequential random	modem CD-ROM
transfer schedule	synchronous asynchronous	tape keyboard
sharing	dedicated sharable	tape keyboard
device speed	latency seek time transfer rate delay between operations	
I/O direction	read only write only read&write	CD-ROM graphics controller disk

Interrupt-Driven I/O Cycle



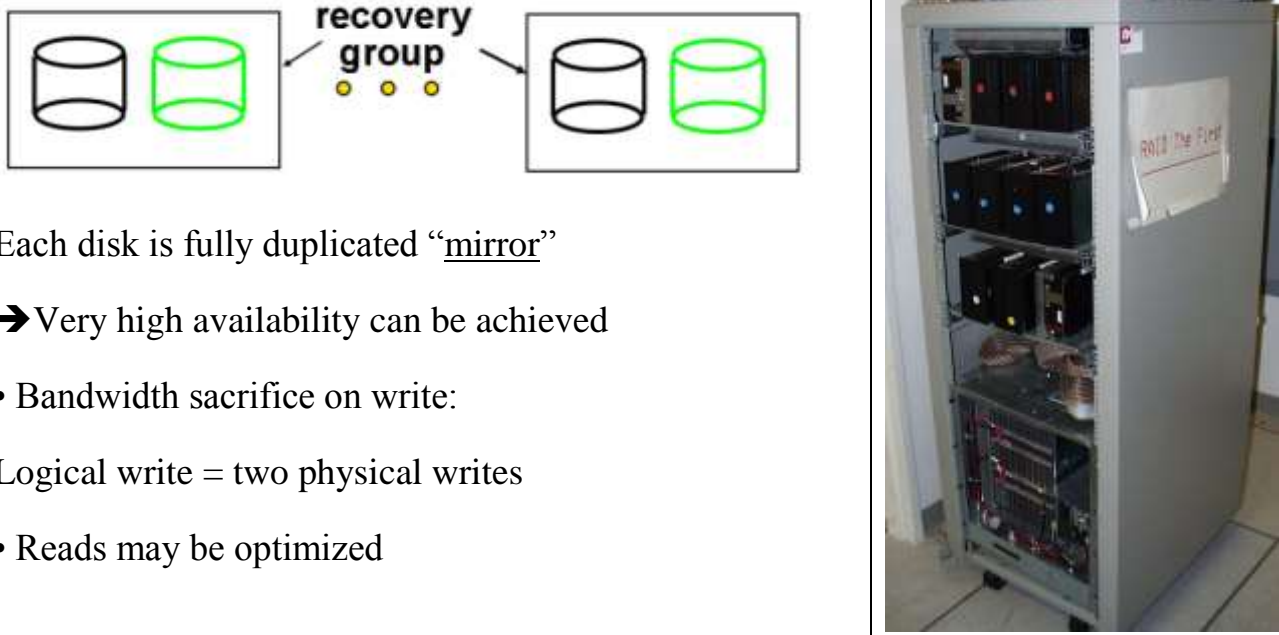
STREAMS

- **STREAM** – a full-duplex communication channel between a user-level process and a device
- A STREAM consists of:
 - **STREAM head** interfaces with the user process
 - **driver end** interfaces with the device
 - zero or more STREAM modules between them.
- Each module contains a **read queue** and a **write queue**
- Message passing is used to communicate between queues



Redundant Arrays of Inexpensive Disks

RAID 1: Disk Mirroring/Shadowing



Each disk is fully duplicated “mirror”

→ Very high availability can be achieved

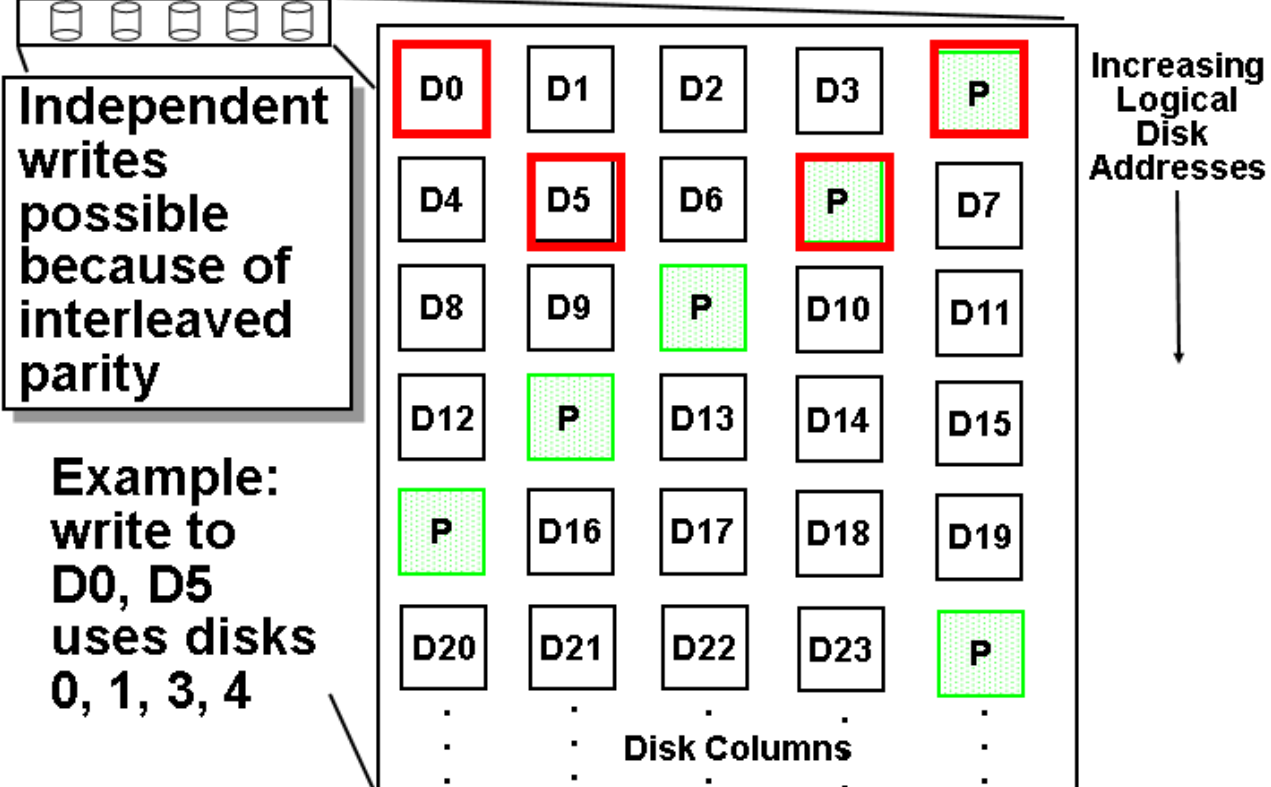
- Bandwidth sacrifice on write:

Logical write = two physical writes

- Reads may be optimized

In **RAID3**, Sum computed across recovery group to protect against hard disk failures, stored in P disk

RAID 5: High I/O Rate Interleaved Parity



Independent writes possible because of interleaved parity

Example: write to D0, D5 uses disks 0, 1, 3, 4

D0	D1	D2	D3	P
D4	D5	D6	P	D7
D8	D9	P	D10	D11
D12	P	D13	D14	D15
P	D16	D17	D18	D19
D20	D21	D22	D23	P
⋮	⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮	⋮

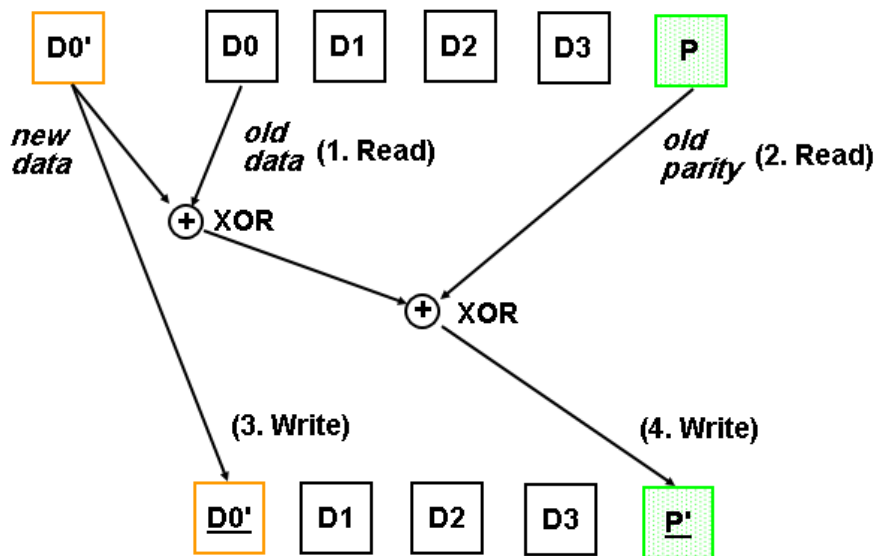
Increasing Logical Disk Addresses

Disk Columns

Problems of Disk Arrays: Small Writes

RAID-5: Small Write Algorithm

1 Logical Write = 2 Physical Reads + 2 Physical Writes



Page 21


Summary: RAID Techniques

- **Disk Mirroring, Shadowing (RAID 1)**

Each disk is fully duplicated onto its "shadow"

Logical write = two physical writes


100% capacity overhead



- **Parity Data Bandwidth Array (RAID 3)**

Parity computed horizontally

Logically a single high data bw disk




- **High I/O Rate Parity Array (RAID 5)**

Interleaved parity blocks

Independent reads and writes

Logical write = 2 reads + 2 writes



Page 22