

11/14

# Storage = SSD & Filesys Basics

Common memory leak from Lab 2.

```
exec() {
```

```
    struct vspace new;
```

```
    // initialize & set up new vspace;
```

- 1 `vspaceinit(&p->vspace);` → what happens to your current page table & data?
  - 2 `vspacecopy(&p->vspace, &new);`
  - 3 `vspaceinstall(p);`
  - 4 `vspacefree(&new);`
- } → this makes a deep copy of the new address vspace, is that what we want?

sysinfo

```
struct vregion {
    enum vr_direction dir; // direction of growth
    uint64_t va_base;      // base of the region
    uint64_t size;         // size of region in bytes
    struct vpi_page *pages; // pointer to array of page_infos
};

struct vspace {
    struct vregion regions[NREGIONS]; // the regions for a process' virtual space
    pml4e_t* pgtbl;                    // process' page table
};
```


→ vspace just stores pointers to things

swap from the old vspace to new one

```
exec() {  
    struct vspace new;  
    // initialize & set up new vspace;  
    struct vspace old = p->vspace; // saves old vspace struct, has  
    p->vspace = new;                // references to old page table &  
    vspaceinstall(p); // switch to new // bookkeeping data  
    vspacefree(&old); // safe to free old vspace.  
}
```

# Solid State Drive


flash memory



disk: sectors (512 byte)

SSD: page (2k-4k) read/write unit

Erase Block (1MB-8MB), erase unit




Erase: erase all pages in the block (set all bits to 1)


single level cell

write: write 0s


(1010010)




Initially (Invalid)



Erase (block) (Erased)



Write (page)  
(page written = Valid, other pages = Erased)



read (page)


# Performance

- 4KB
- read (page level) : good (~10µs) for sequential and random read
  - write (page level) : write to an erased page (~100µs) / program
  - erase (block level) : takes a few ms

\* can only program/erase so many times before the block becomes unusable

→ wear leveling - spread writes across blocks & pages

On erase, what happens to existing data in the block?



to overwrite page 3 we need to erase block 1

→ copy existing data to erased pages elsewhere, erase block and then write to block 1 pages

not the most efficient approach, had to move data around multiple times

## Flash Translation Layer

→ logical to physical block #  
↓  
what we give to client (kernel)

Block 1 → ~~Block 1~~  
          ↓  
          Block 2.

(page #, block not exposed to client)

# Performance Metrics

- I/O Operations per Second (IOPS)
- I/O request: size (# of bytes to read/write)
- access patterns matter! (workload)

→ simple average IOPS

$$\frac{1}{\text{average request latency}}$$

for disk:  $\frac{1}{14 \text{ ms}} \approx 70 \text{ IOPS}$

→ actual IOPS

$$\frac{\text{\# of I/O requests}}{\text{total time to complete the requests}}$$

for 10 sequential reads on disk


$$\frac{10}{14 \text{ ms}} \approx 700 \text{ IOPS}$$

980 PRO PCIe® 4.0 NVMe™ SSD 1TB

Total \$109.99 ~~\$159.99~~

[BENEFITS](#) [SPECS](#) [REVIEWS](#) [SUPPORT](#) [RELATED](#)

[Chat about Black Friday Deals](#)

 Save an additional 5% with 3 pack!



Sequential Reads  
up to **7,000 MB/s**

Sequential Writes  
up to **5,100 MB/s**

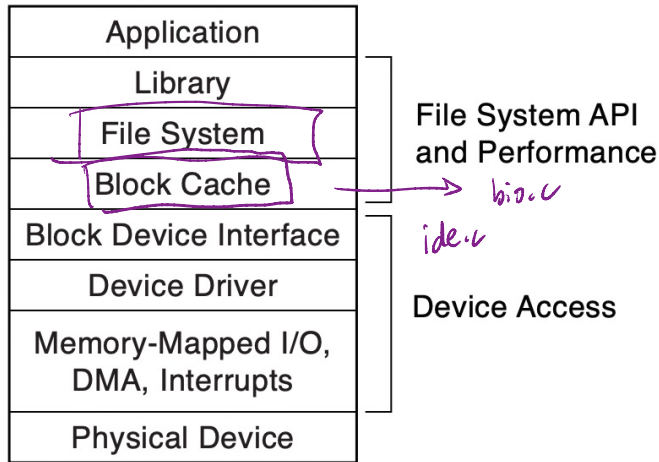


Random Reads  
up to **1,000K IOPS**

Random Writes  
up to **1,000K IOPS**

# Filesystems

- abstraction for storage devices
- names & data (files, directories)
- manages storage devices & offers richer APIs



Data vs Metadata

↓  
file data

↓  
info about file  
(size, owner, permission ...)