Read Disturb Management Improvement in SSDs

Qingru Meng

Intel Corp.

Flash Memory Summit

Outline

- Read Disturb in NAND Flash Array
- Read Disturb Management in SSDs
  - Read Disturb Impact Metrics
  - Read Disturb Management by Different Vendors/Products
  - Improvement of Read Disturb Management
- Summary

Read Disturb in NAND Flash Array

- High V_{pass} causes unwanted charge gain
- Happens to deselected WLs
- Electric-field driven, so mostly L0
- Worse with cycling
- More noticeable in TLC and QLC

Read Disturb must be dealt with to prevent data corruption

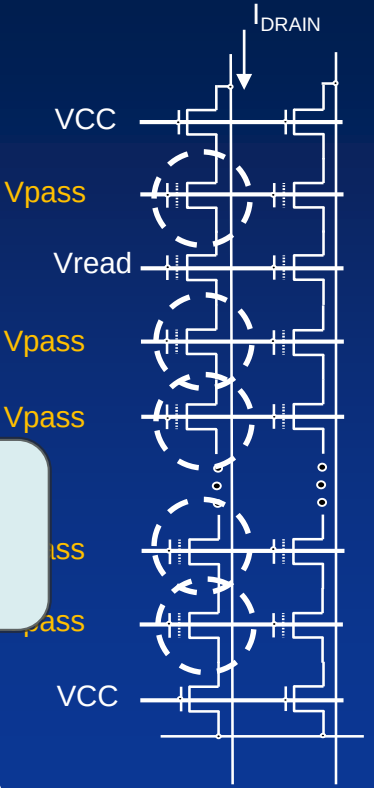

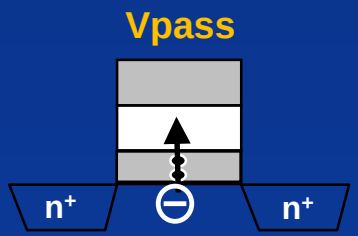

Flash Memory Summit

Outline

- Read Disturb in Flash Memory Array
- Read Disturb Management in SSDs
  - Read Disturb Impact Metrics
  - Read Disturb Management by different vendors/products
  - Improvement of Read Disturb Management
- Summary

Read Disturb Management

- Objective: refresh the data before read disturb causes ECC failures
- Tasks:
  - Track/Count the reads
  - Make room for data relocation
  - Copy data to a new location
  - Update LBA map and invalidate old data

All sorts of SSD internal activities run in parallel with host data requests ➔ Performance Impact

Flash Memory Summit

Outline

- Read Disturb in Flash Memory Array
- Read Disturb Management in SSDs
  - Read Disturb Impact Metrics
  - Read Disturb Management by different vendors/products
  - Improvement of Read Disturb Management
- Summary

Read Disturb Impact Metrics

- IOPs/IO Stability Drop
- Latency/QoS Increase
- IOPs Recovery Time
- Triggering Rate/Eviction Rate
- Write/Read Ratio in 100% read workloads

For all these metrics, smaller is better

Example (I)

- 30% IOPs Drop
- Takes more than half hour to recover
- Drop Occurs every 180 drv reads
- Average/Max latency increases more than 40%

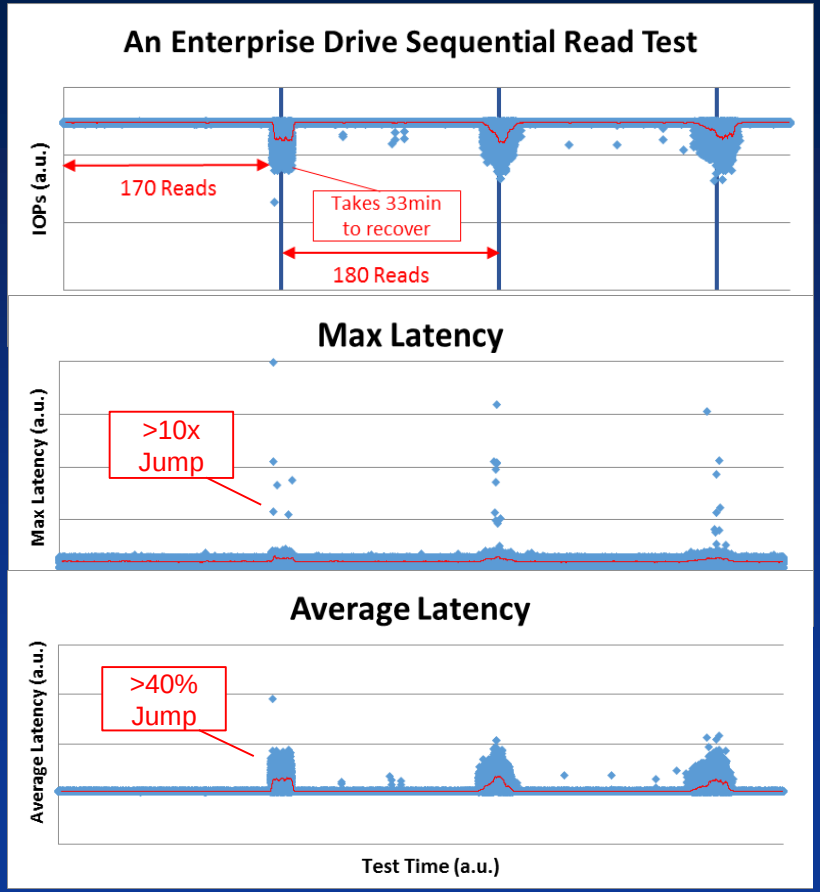

Example (II)

- NAND Write/Host Read ratio:
  - 1e-3 to 1e-4.
  - Random read: triggering earlier and higher write/read ratio
- Power varies when data relocation takes place

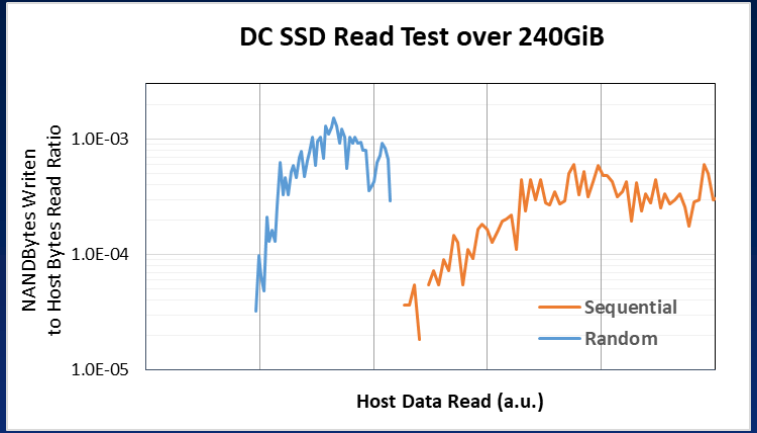

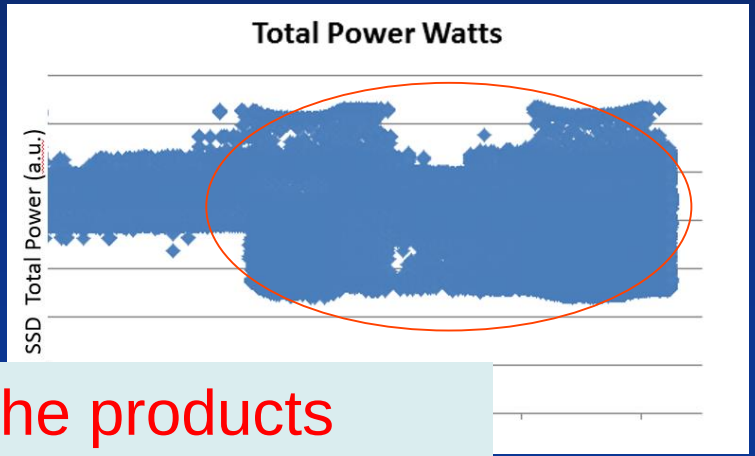

Not all signatures manifest in all the products
Improvement➔ Less visible in certain metrics

Flash Memory Summit

Outline

- Read Disturb in Flash Memory Array
- Read Disturb Management in SSDs
  - Read Disturb Impact Metrics
  - Read Disturb Management by different vendors/products
  - Improvement of Read Disturb Management over Time
- Summary

Read Disturb Handling by Different Vendors (Client SSDs)

All 64L TLC NAND Based NVMe SSDs

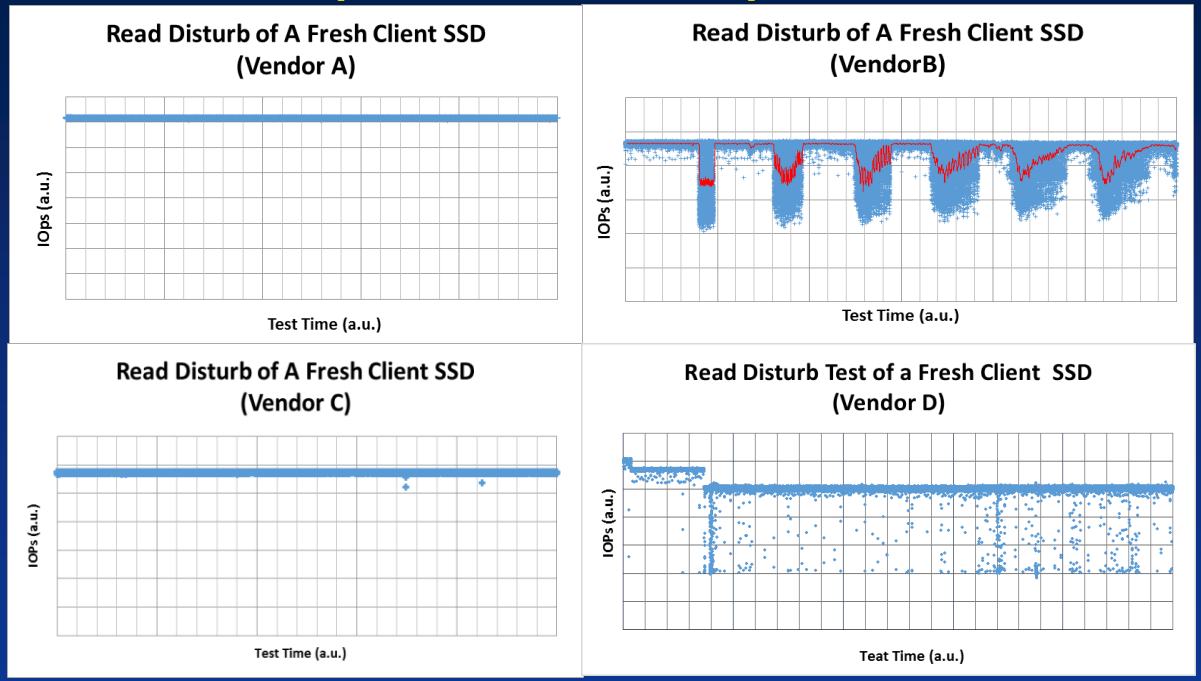

Some handles better than others

Read Disturb Management by Different Vendors (DC/Enterprise)

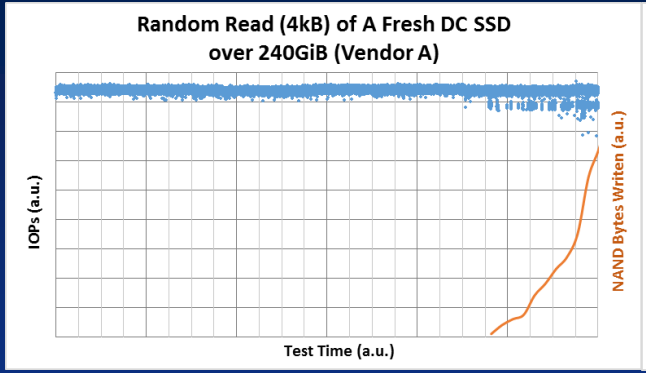

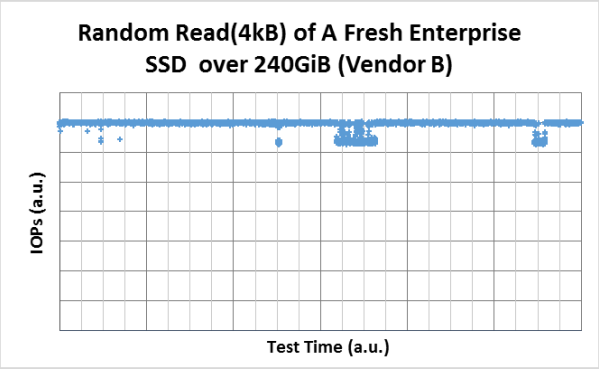

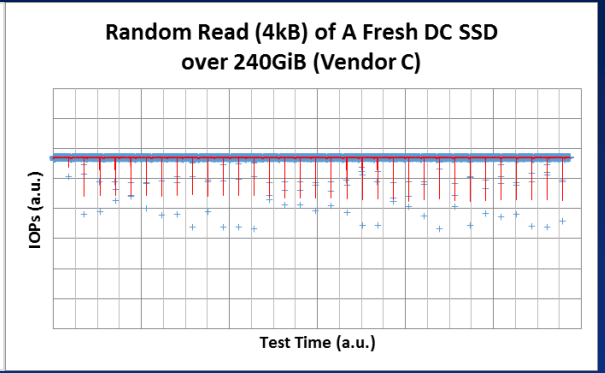

Two methods:

- Deal with it only when you have to
- Handle it regularly

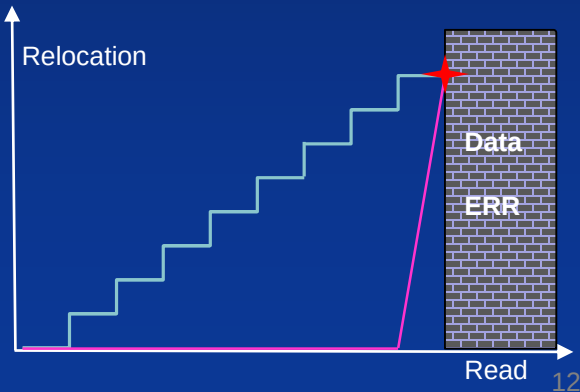

Read Disturb induced ECC Failures

- Delay/Inadequate management can lead to ECC failures

Not expected in today's DC SSDs

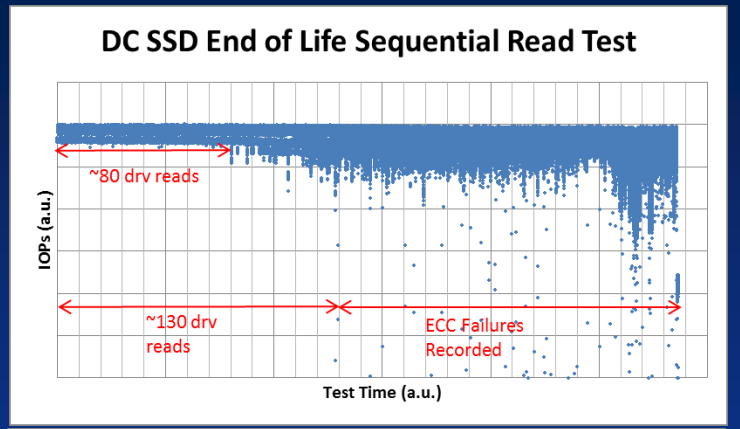

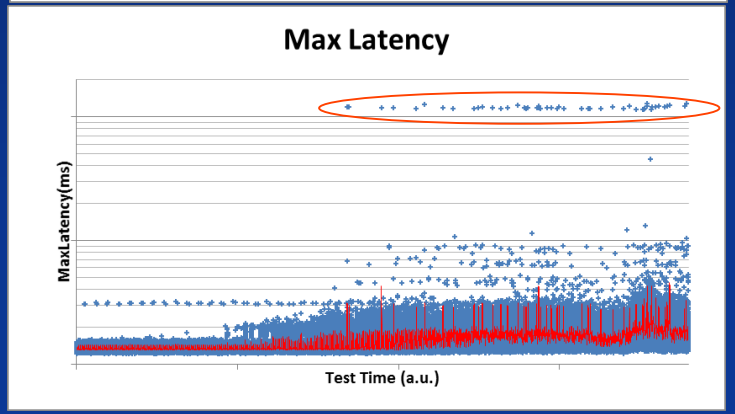

Outline

- Read Disturb in Flash Memory Array
- Read Disturb in SSDs
  - Read Disturb Impact Metrics
  - Read Disturb Management by different vendors/products
- Improvement of Read Disturb Management
- Summary

Improvement in Client SSDs

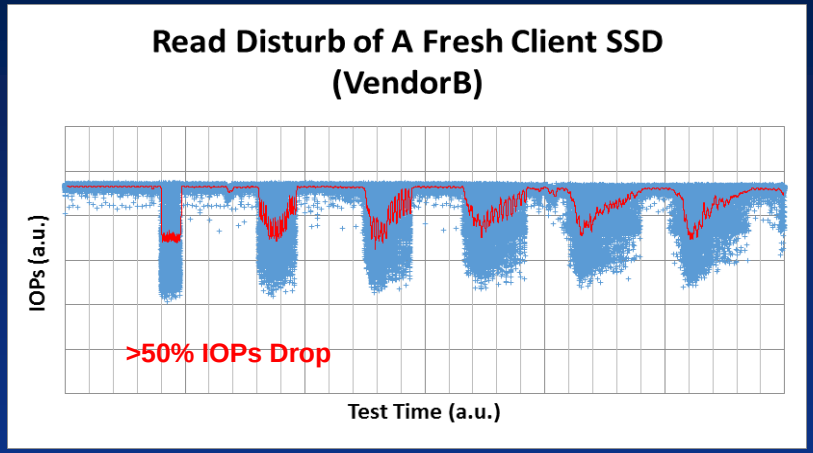

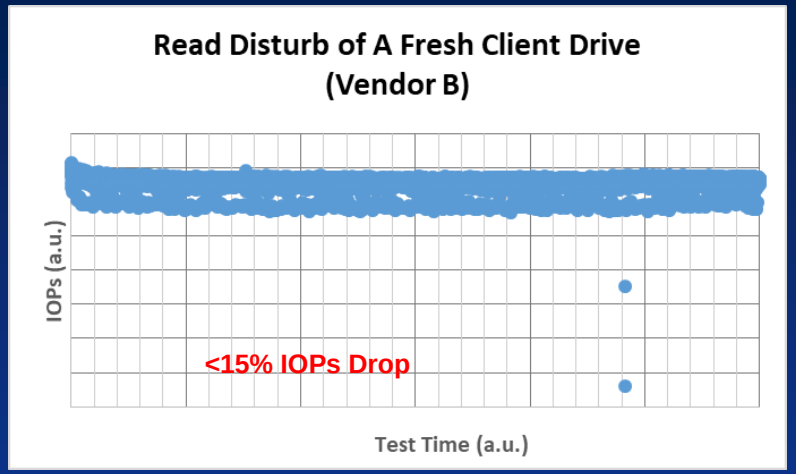

New SSD with same controller, different algorithm and much reduced IOPs drop

Improvement in DC SSDs

Contributing Factors:

- Better NAND (lower RBER and faster read/write)
- Richer features (erase/program/data transfer suspend)
- More powerful ASIC (processing power and ECC engine)
- Smarter read disturb policy

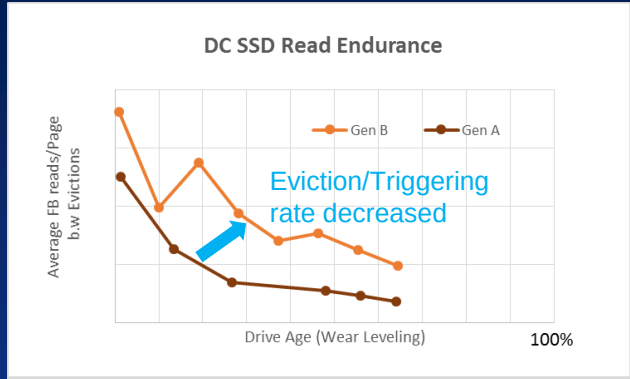

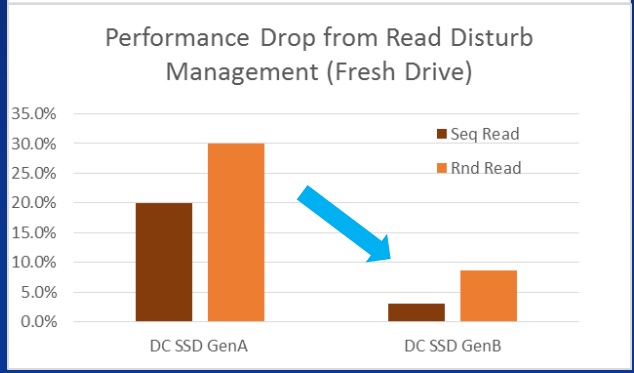

Summary

- Read disturb management is critical to a drive's performance and reliability
- Industry has dramatically improved read disturb management over the years
- Thorough validation and assessment of read disturb policy are key to a robust SSD.

THANK YOU!