STORAGE DEVELOPER CONFERENCE
SDC 22
Fremont, CA
September 12-15, 2022
BY Developers FOR Developers

A SNIA Event

An Approach for Impact Analysis of Flash Behavior on QoS in DC/Enterprise SSDs

Ravishankar Singh ,Senior Staff Engineer
Yogesh Khurana ,Staff Engineer
SSIR

Agenda

- Introduction
- NAND Basics
- SSD Architecture
  - IO Operations
  - Garbage Collection
  - Over-provisioning
- QoS Parameters
- QoS Metrics
- Factors for QoS Behaviour
- NAND parameters impact on QoS
- SW Features impact on QoS
- Results
- References

Introduction

Introduction

- Performance is key feature of SSD Specification for enterprise customers
- SSDs performance are always evolving because of NAND speed and host interface speed enhancements.
- It is very challenging to maintain Quality of Service(QoS) with enterprise/DC workloads.
- NAND parameters and FW behaviors have impact on QoS.

NAND Basics

- NAND Flash is logically divided into blocks and blocks consists pages
  - Basic Erase Operations supported by block unit
  - Program/Read operations are supported by Page unit

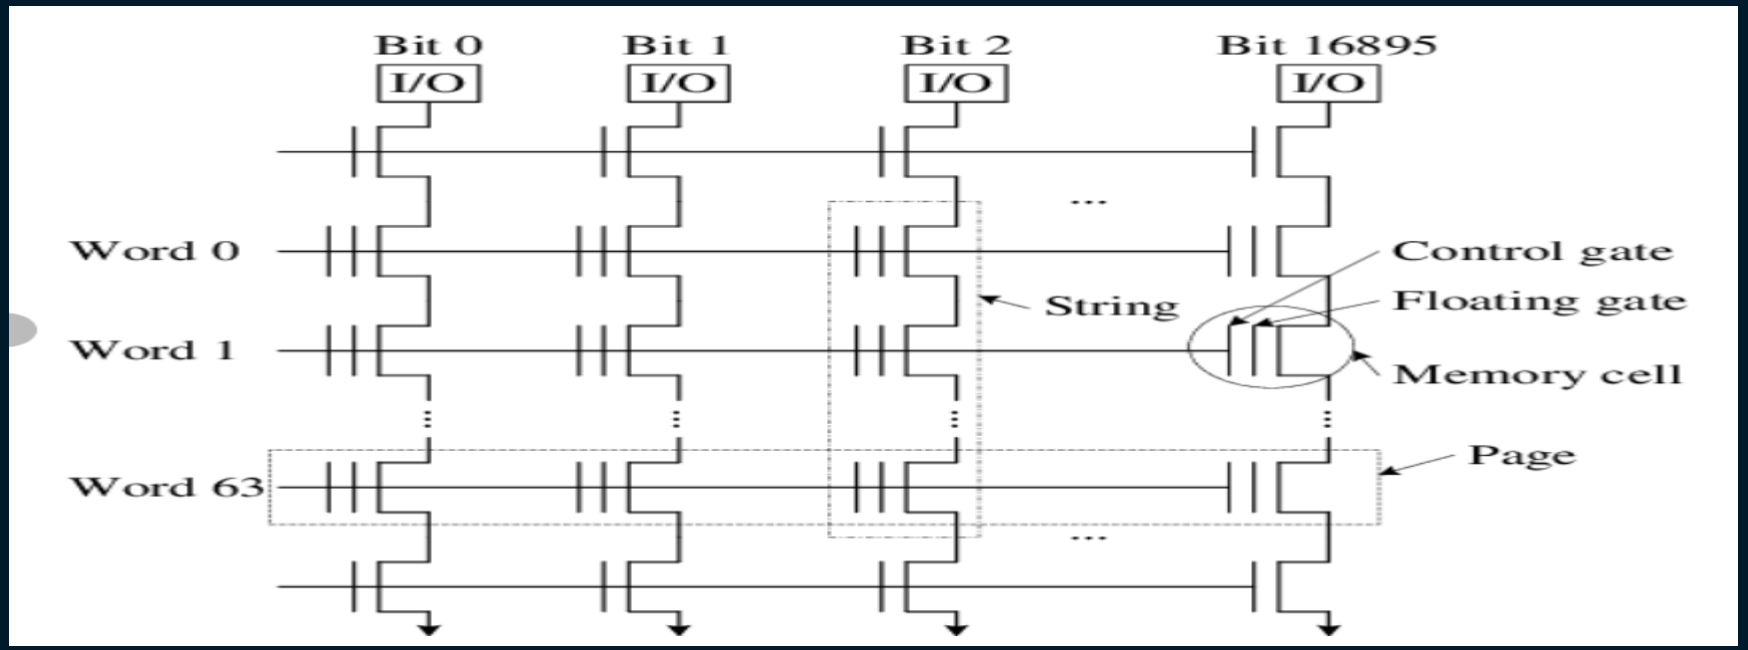

Reference :-https://www.researchgate.net/figure/Cell-Layout-of-NAND-Flash-Memory_fig1_234126811

NAND Basics

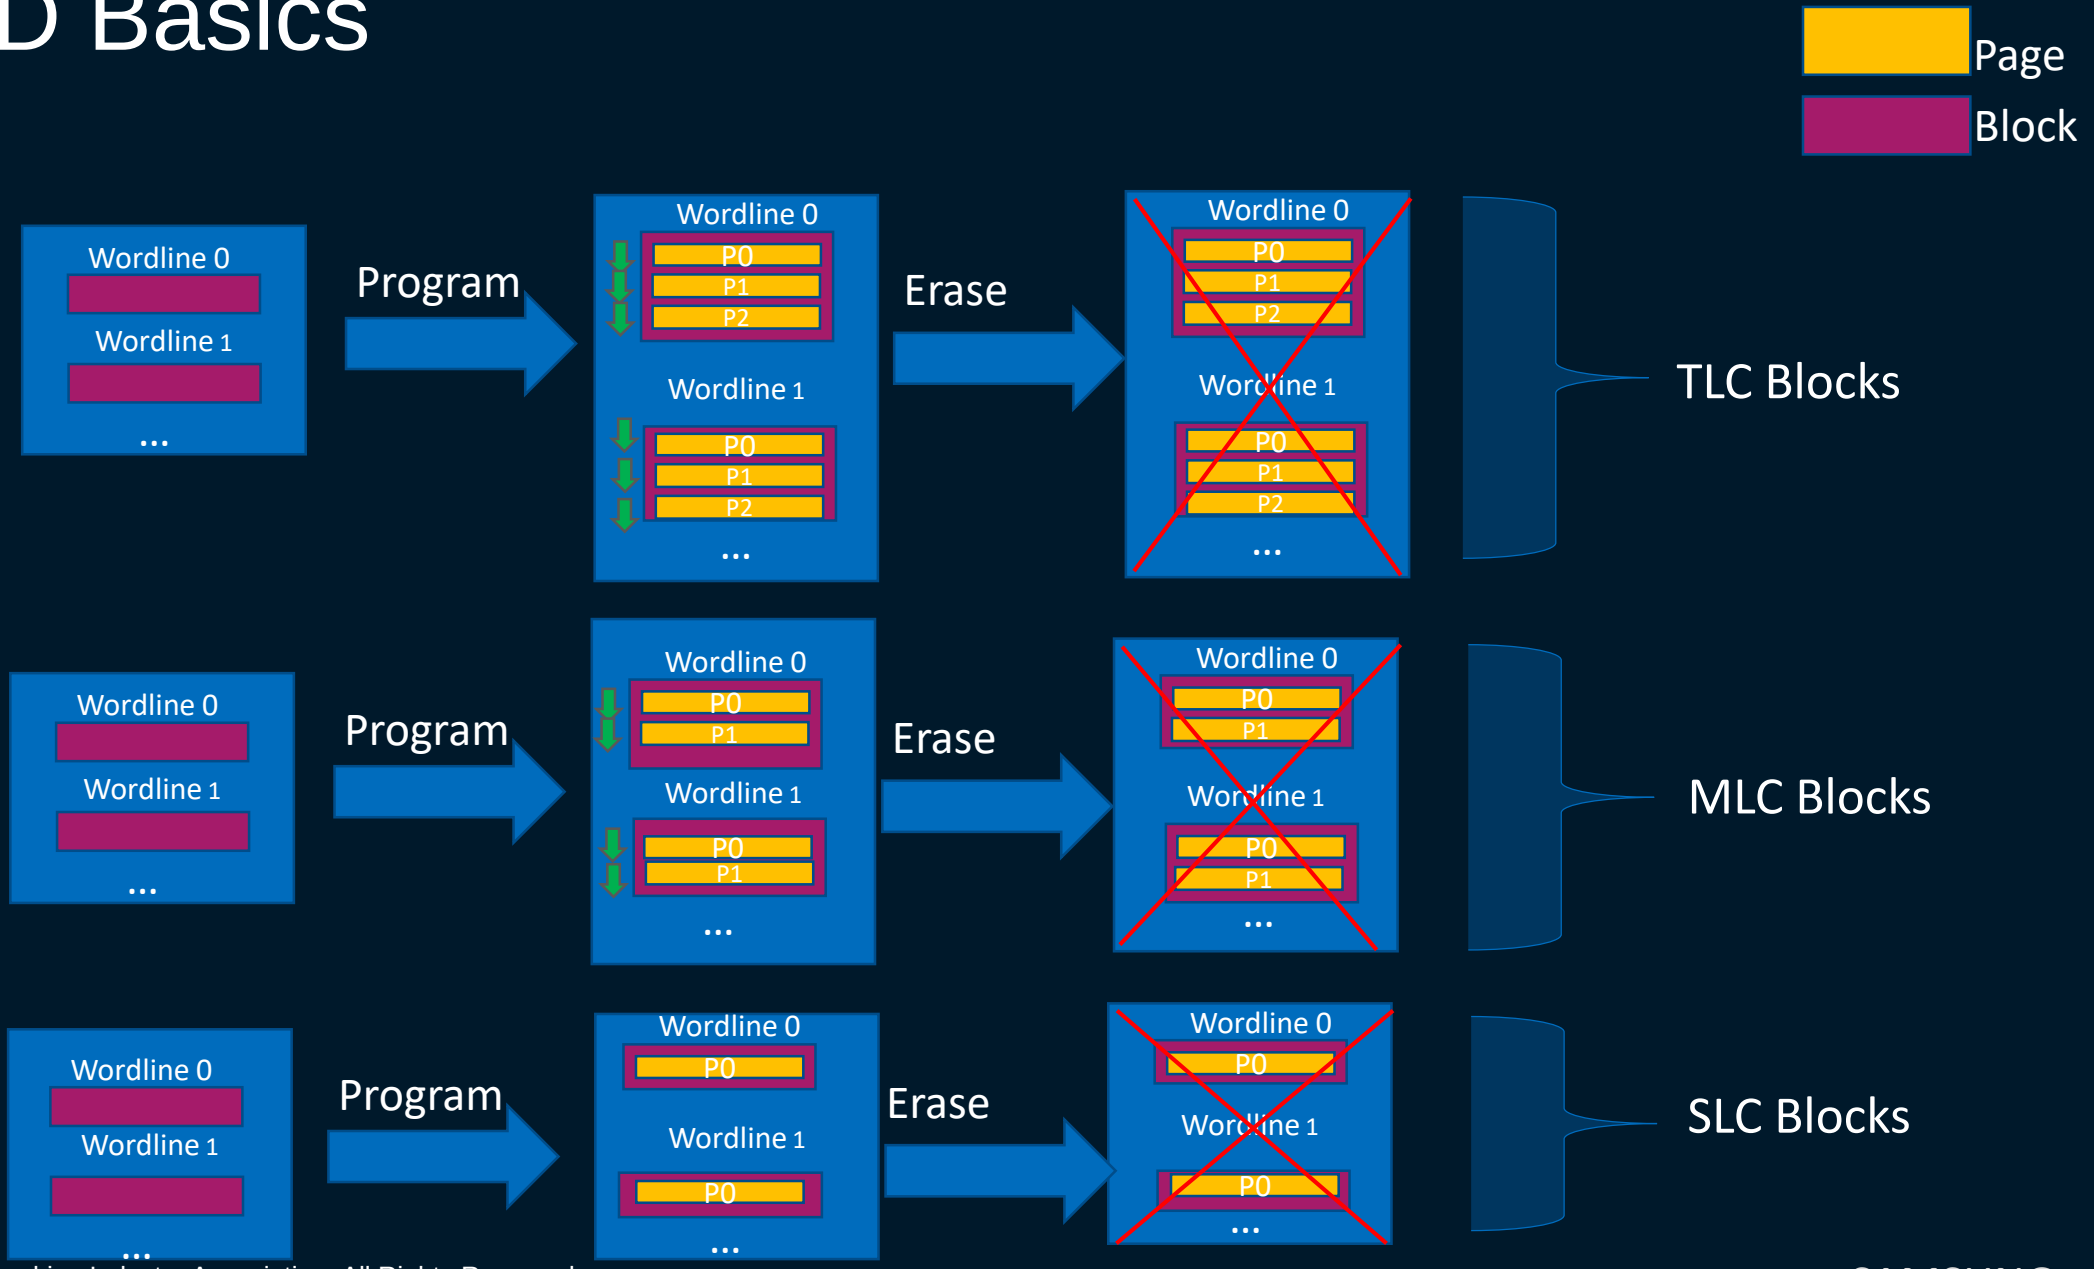

SAMSUNG

SSD Architecture

SSD Overall Geometry

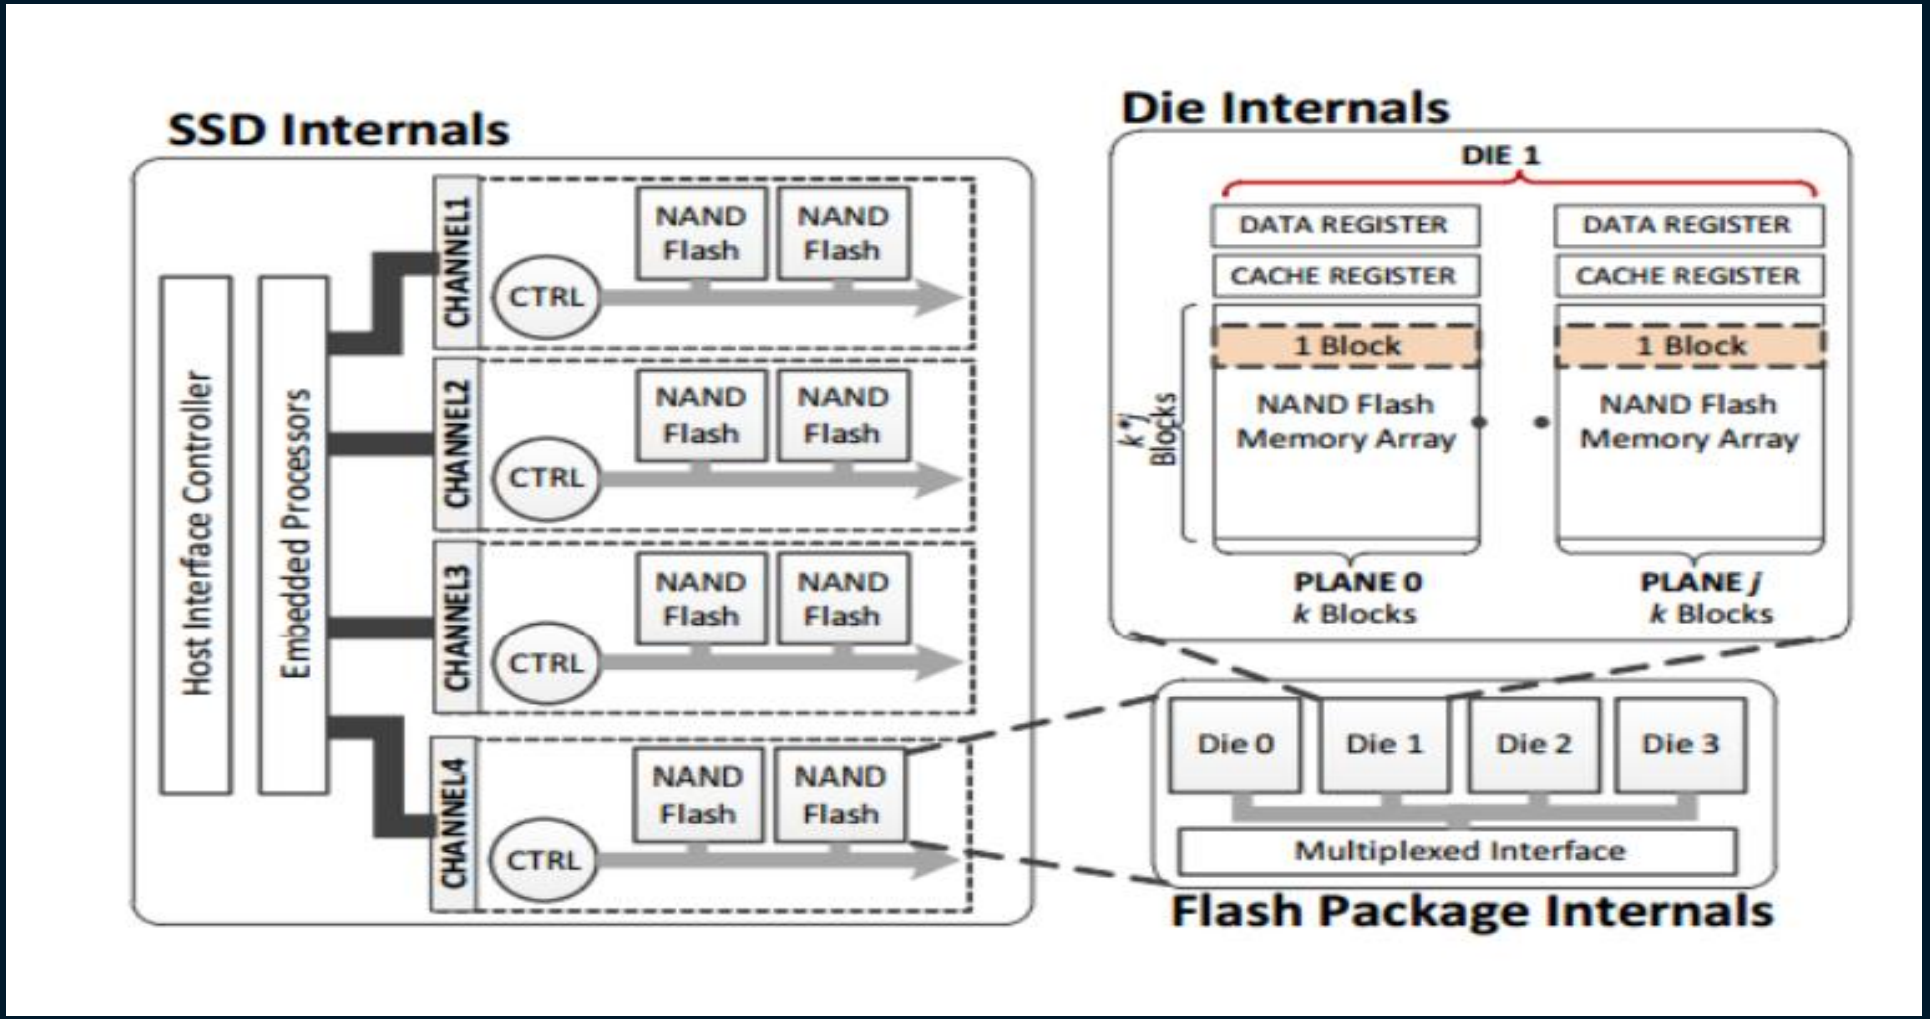

Reference :- https://www.researchgate.net/figure/Physical-internal-architecture-of-SSD_fig1_241633959

SSD IO Operation (Write)

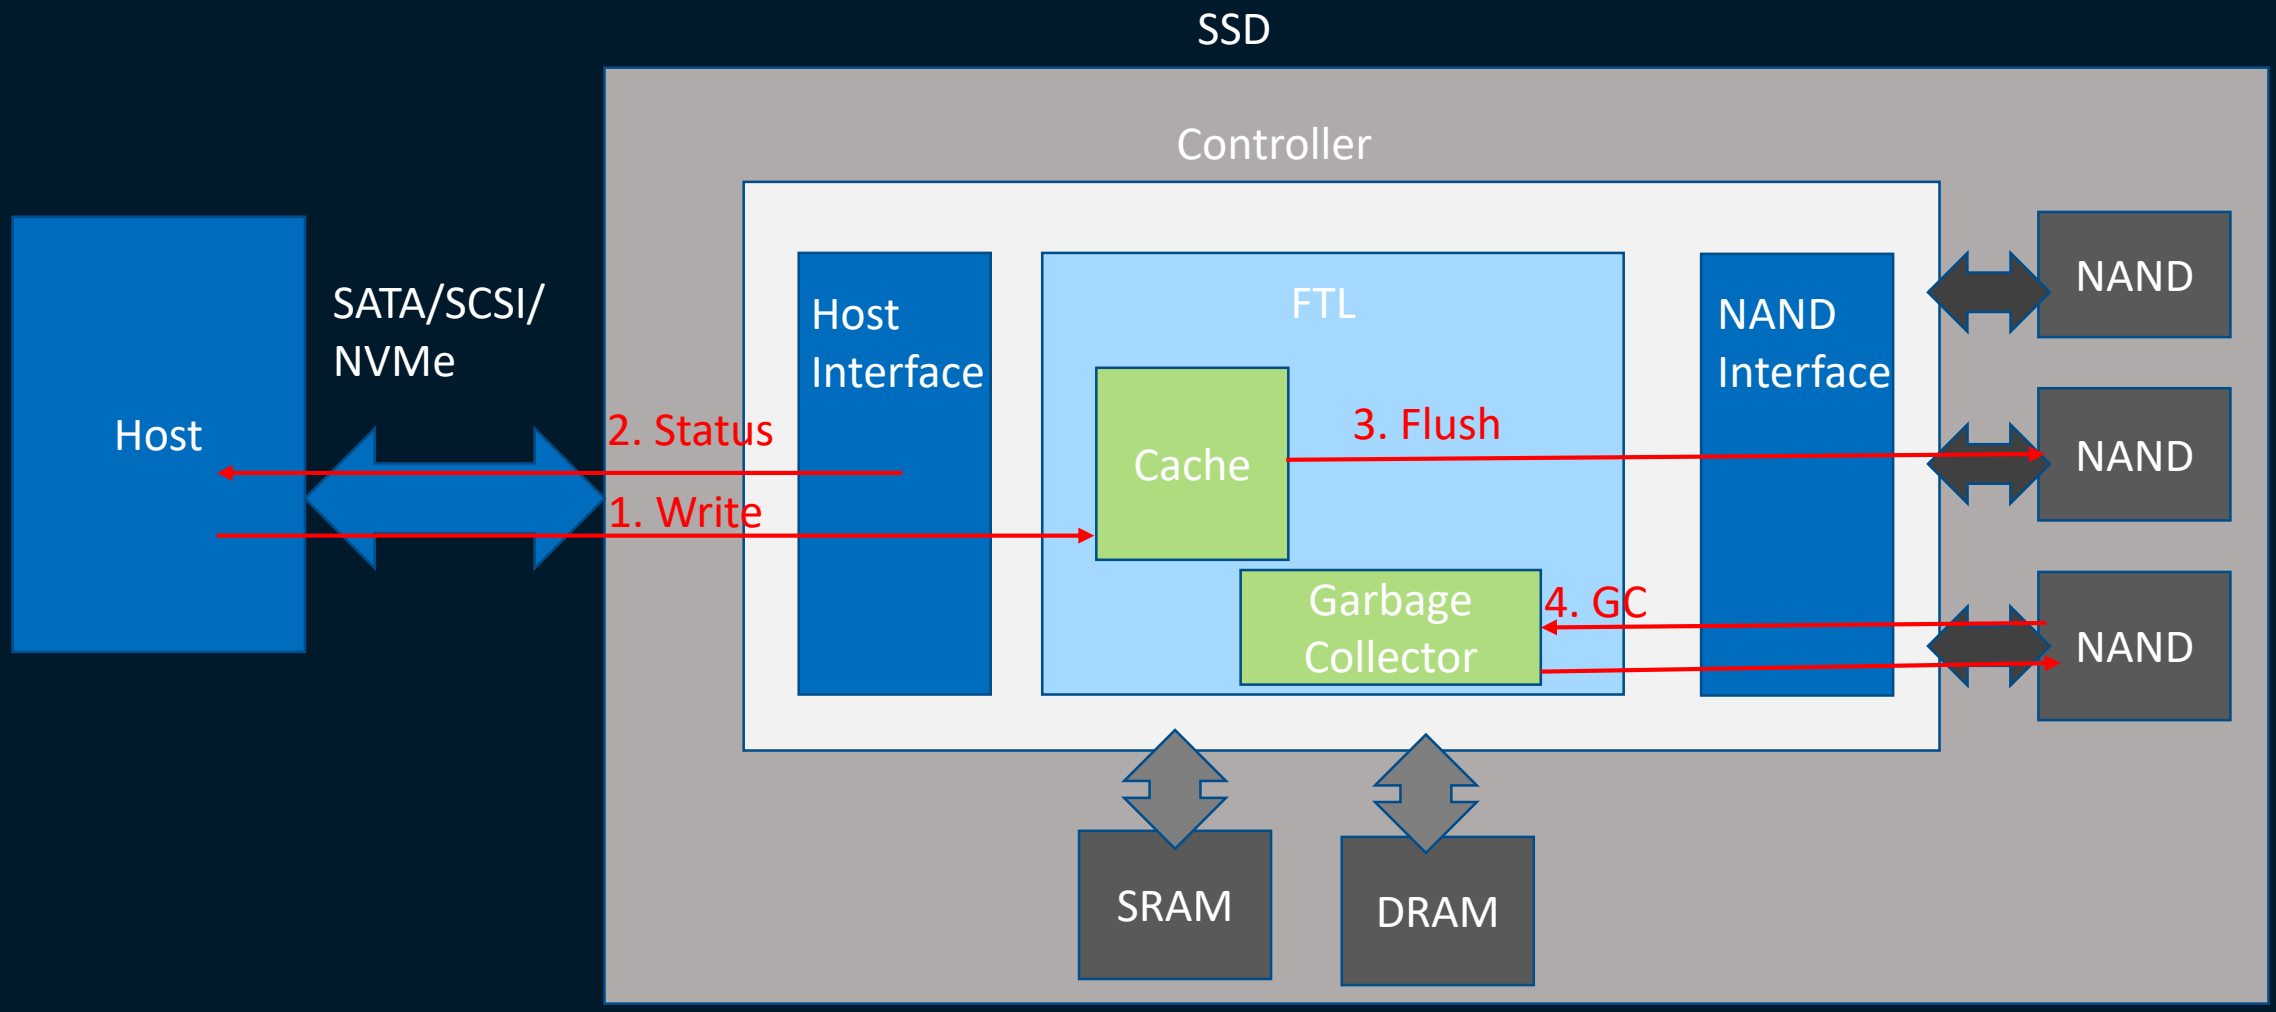

SAMSUNG

SSD IO Operation (Read)

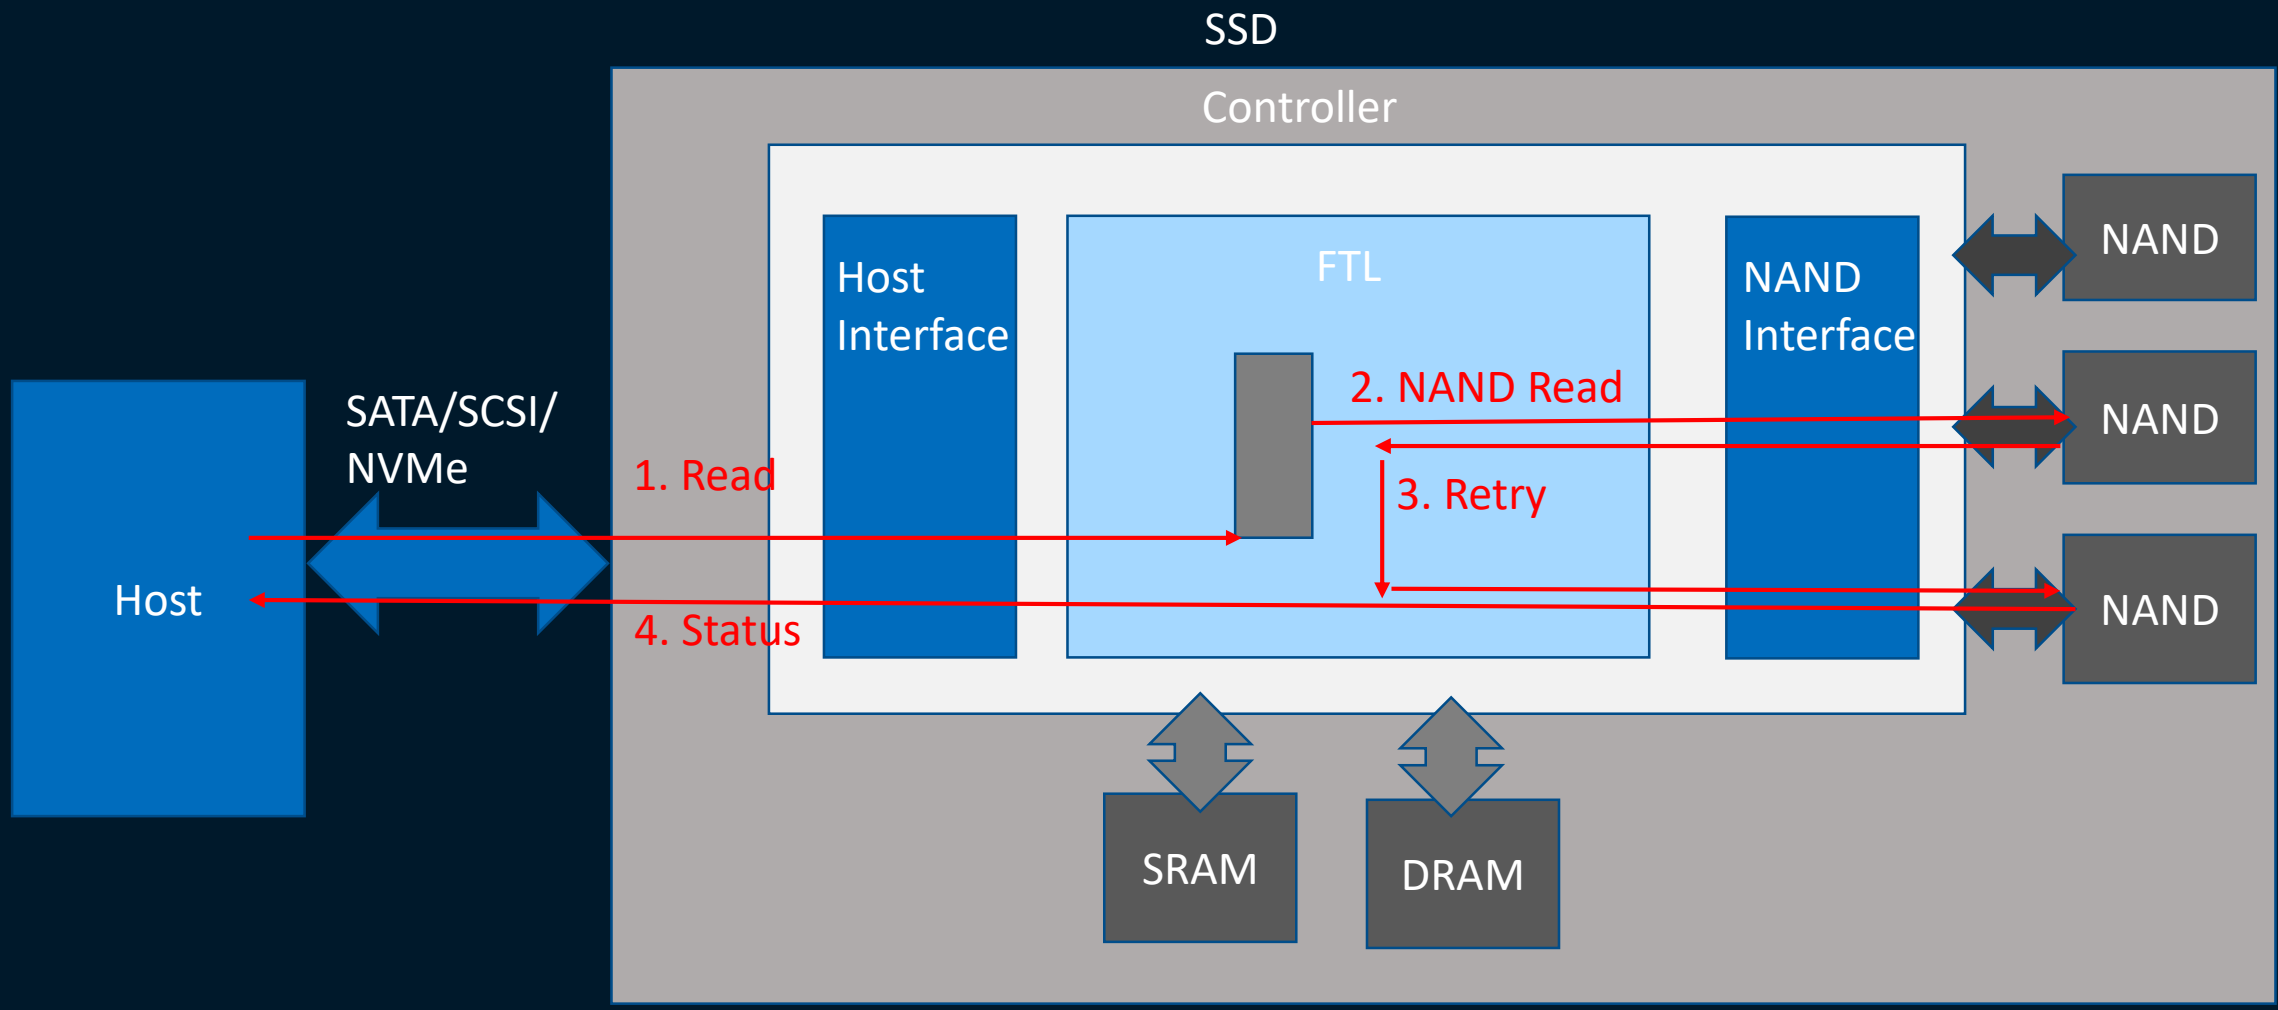

SAMSUNG

Garbage Collection

- How
  - Identifies Which page contain stale data
  - Moves the pages with good data to another block
  - Erase all the data from original block

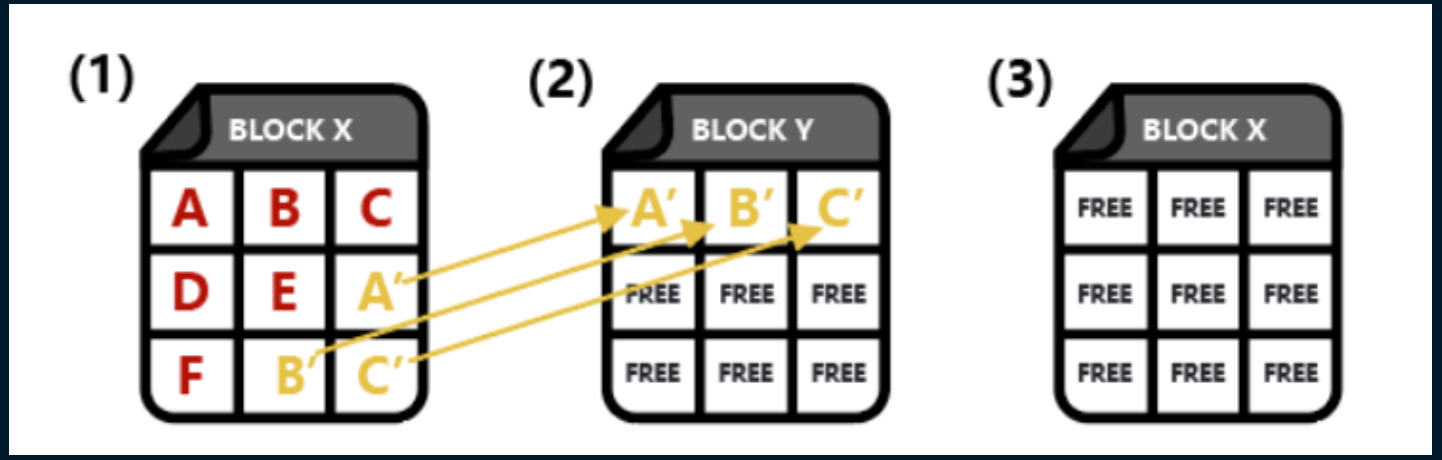

- Why
  - To Manage P/E cycle
  - To reduce impact on Endurance
  - To reduce impact on performance

Reference:-https://ssstc.com/industrial-ssd-features/garbage-collection-ssd/

Over-Provisioning

- How to Calculate
  - Every SSD has fixed NAND chips size
  - User capacity is provide by IDEMA standard.
  - SSD Over-provisioning defined by
    - % OP = (NAND Flash Size – IDEMA Size)*100 / NAND Flash Size

- Why (Impact)
  - High OP  = Low GC  = Low WAF  = High Write Performance 

Quality Of Service (QoS)

- What
  - QoS is quality level of steady and consistent performance for all requested processes
  - It helps to ensure that a particular workload always gets a certain performance level.
  - For better QoS, All processes must finish within specific time limit or above a target confidence value.
- Why is QoS a challenge?
  - SSD policies and characteristics interferes into IO completions and latencies are high

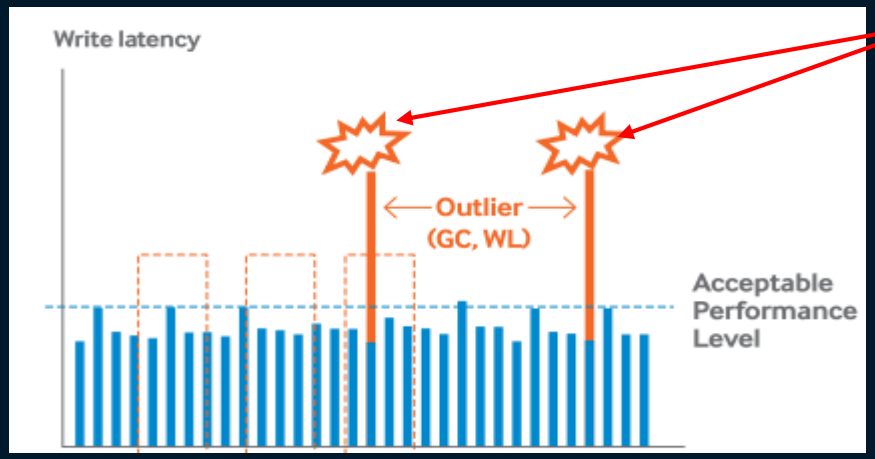

Very High latency because of GC or WL.
QoS is compromised

Reference:-http://www.samsung.com/global/business/semiconductor/minisite/SSD/downloads/document/Samsung_SSD_845DC_06_Quality_of_Service(QoS).pdf

QoS Parameters

- Workload Access
  - Read
  - Write
  - Mixed
  - Trim
- Queue Depth (Based on Device type)
  - 32 for SATA / 64K for NVMe
- Block Size
  - 128KB (Sequential)
  - 4KB Random
- Device State
  - Clean
  - Sustained

QoS Metrics

- QoS Latency
  - Average latency
  - Lower Nines : 99%, 99.9%, 99.99%
  - Higher Nines : 99.999%, 99.9999%, 99.99999%, 99.999999%
- Consistency (Confidence Value)
  - More than 95%

Factors for QoS Behaviour

- Host Factors
  - CPU mode and Speed
  - Schedulers (NOOP, CFQ)
  - RAM Speed
  - Host Protocol interface (PCIe, SATA, SAS)
- Device Factors
  - NAND Parameters
    - NAND timings (tR/tPROG/tBERS)
    - NAND Interface Speed
    - NAND Geometry(Channel/Ways)
  - FW Features
    - OP
    - PLP
    - Program / Erase suspend and resume
    - GC

NAND Parameters impacts on QoS

- NAND timings
  - tR – Page read time
  - tPROG – SLC/TLC/MLC page program time
  - tBERS – Block erase time
  - Better tR/tPROG/tBERS timings means improved latency
  - tR/tPROG has impact on program/erase suspend and resume latency
- NAND Geometry
  - High interleaving with better channel / ways geometry.
  - More Outstanding commands are processed parallel
  - Improved QoS latency with high parallelism
- NAND Interface speed
  - Better NAND interface ,means improved latency

FW Features impact on QoS

- PLP (Power loss Protection)
  - Safegaurds SSD in case of Power-loss
  - In case of power failure, PLP helps to flush in-flight data to flash
  - Cache will be turned off, hence performance degradation
- Over-provisioning
  - High OP provides low GC overhead hence better IO latencies
  - OP has impact on WAF (NAND Write vs Host write)
    - Read Intensive SSDs has 10% OP and high WAF
    - Mixed Pattern SSDs has 30% OP and Low WAF

FW Features (Cont…)

- Suspend / resume operations
  - Read IO has higher priority than Program/Erase
  - Program/erase suspended if it interrupts with Read IO
  - High IO latency because of suspend / resume operation
  - Mixed workloads has high impact

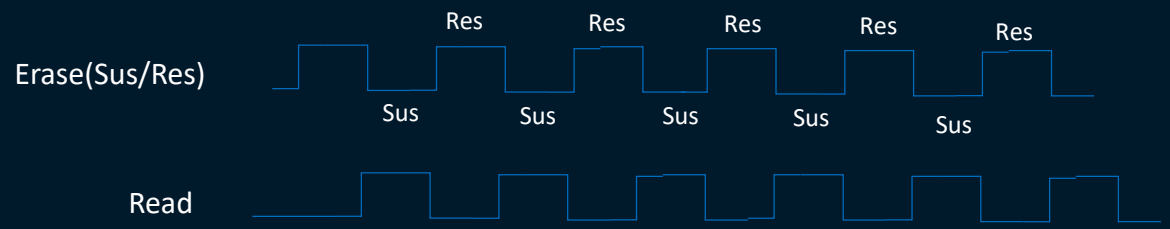

  - IO operations are delayed because of $t_{SUS/RES}$ + $t_{RES/SUS}$ latencies. Hence QoS drop will be observed.

SAMSUNG

Results

- QoS bottleneck were analyzed across multiple drives
- Below Table extracts the QoS results with flash parameter changes
- All the results are relative and no absolute value will be shared

Parameter Changes	Read QoS	Write QoS	Mixed QoS	Reason for QoS Drop
tR Drop : 10%	26% ↓	12% ↓	20% ↓	Slow NAND Read
Die Size Changes (2*256Gb->1*512Gb)	44% ↓	32% ↓	39% ↓	Reduced Interleaving
NAND Type changes(MLC->TLC)	25% ↓	65% ↑	30% ↑	Better bit Density improves Program latency

References

- https://www.researchgate.net/figure/Physical-internal-architecture-of-SSD_fig1_241633959
- https://www.researchgate.net/figure/Cell-Layout-of-NAND-Flash-Memory_fig1_234126811
- https://ssstc.com/industrial-ssd-features/garbage-collection-ssd/
- http://www.samsung.com/global/business/semiconductor/minisite/SSD/downloads/document/Samsung_SSD_845DC_06_Quality_of_Service(QoS).pdf
- https://www.snia.org/sites/default/files/SDC/2019/presentations/backups/Konan_Andrei_Origin_of_million_IOPs_throughput_and_Micro_Second_Latencies_in_NVMe_Enterprise_SSDs.pdf

Acknowledgements

- Puneet Kukreja
- Mihir Shishukala
- Mohan Raj Veerasamy
- Amit Devgan
- Jaegyu Choi
- Madan Udaykumar

Thank You !

Our Contacts

ravi.s@samsung.com

y.khurana@samsung.com

SAMSUNG

Please take a moment to rate this session.

Your feedback is important to us.

