



Flash Memory Summit


QLC Challenges

QLC SSD's Require Deep FTL Tuning
Karl Schuh – Micron



Flash Memory Summit

The Wonders of QLC



- Cost
- Capacity
- Performance
 - Error Rate depends upon compensation for transaction history
- Endurance
 - Data Organization is Key
 - Complex NAND Management algorithms required to support the target Total Bytes Written

Not 3 bits-per-cell to 4 bits-per-cell, but 8 charge levels to **16 charge levels!**



Flash Memory Summit


Performance and Endurance

Inextricably interconnected in QLC

- Data Organization and Steerage – driven by the workload
 - SLC, MLC and TLC caching to filter high frequency data, and smooth throughput
 - Read / Write mix
 - Logical Saturation
 - Over-Write and Deallocation frequency
 - Steams aid in Data Separation
- Write Throughput
 - Careful attention to history and environment
 - NAND characterization
- Read Latency
 - Continuous Tuning
 - Complex algorithms based on transaction history and characterization

Data Organization

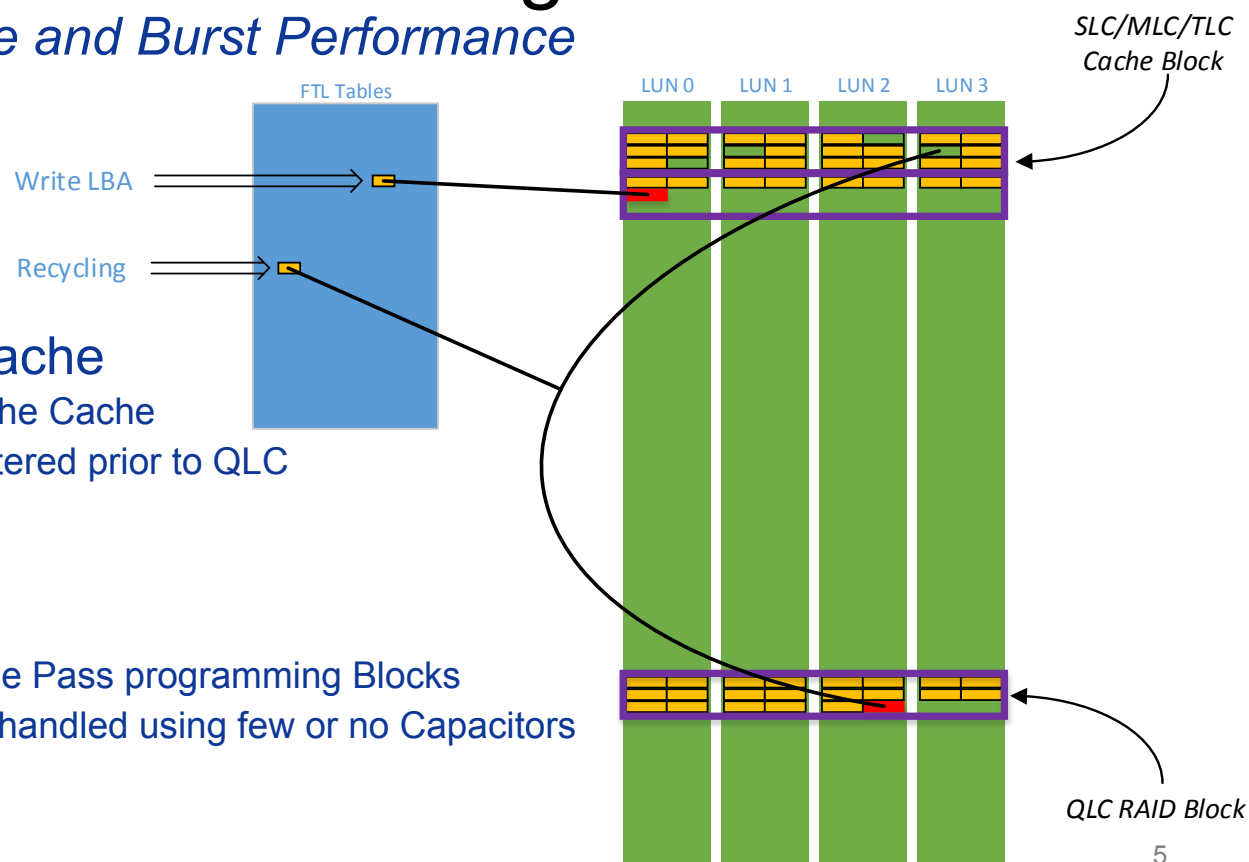
Traditional



- Flat FTL tables
 - Direct lookup for Reads
 - Writes mapped sequentially to a RAID Block across all die
 - Maximize active die, which in turn maximizes throughput
- Hierarchical tables are functionally equivalent

Intermediate Caching

Endurance and Burst Performance

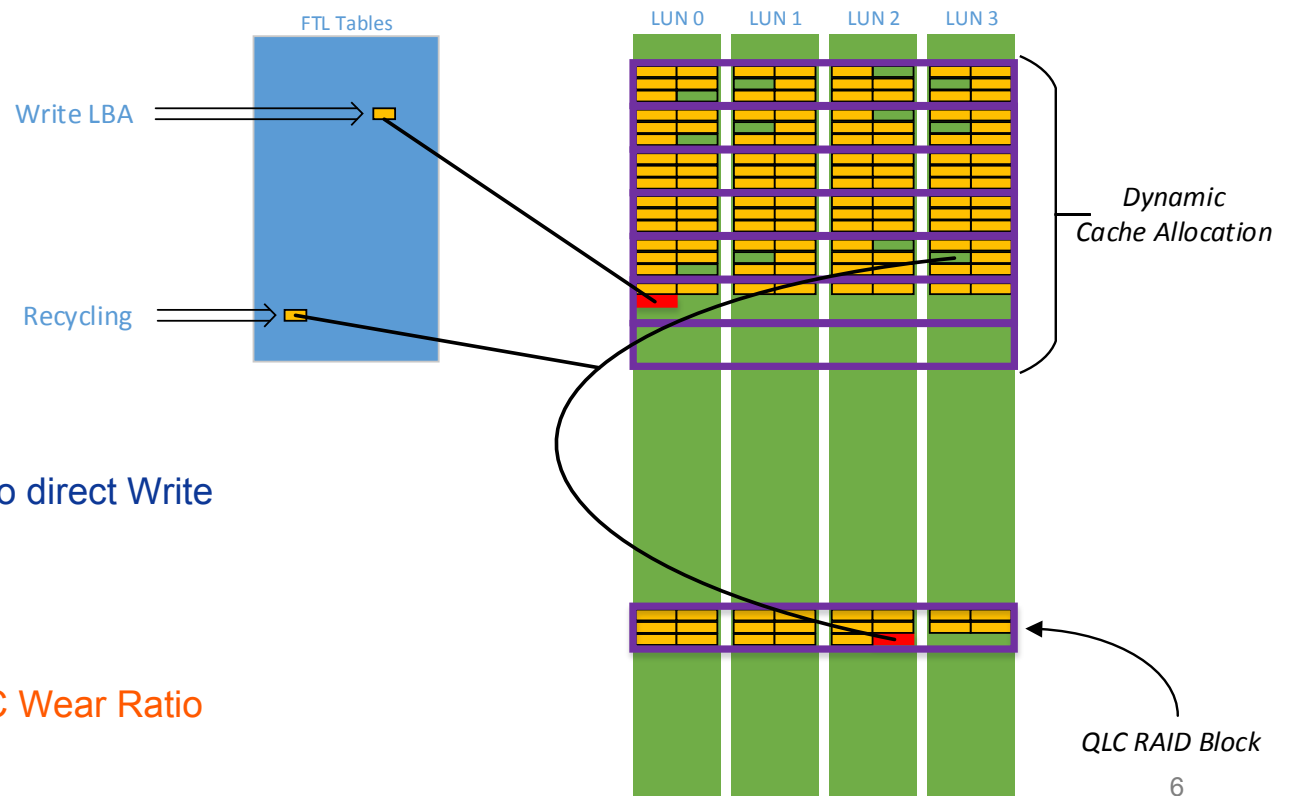


- **SLC or MLC or TLC cache**
 - Write bursts absorbed by the Cache
 - High frequency updates filtered prior to QLC
 - Low FUA latency
 - Improved Endurance
- **Reliability**
 - Host Writes stored in Single Pass programming Blocks
 - Asynchronous power loss handled using few or no Capacitors

Dynamic Caching

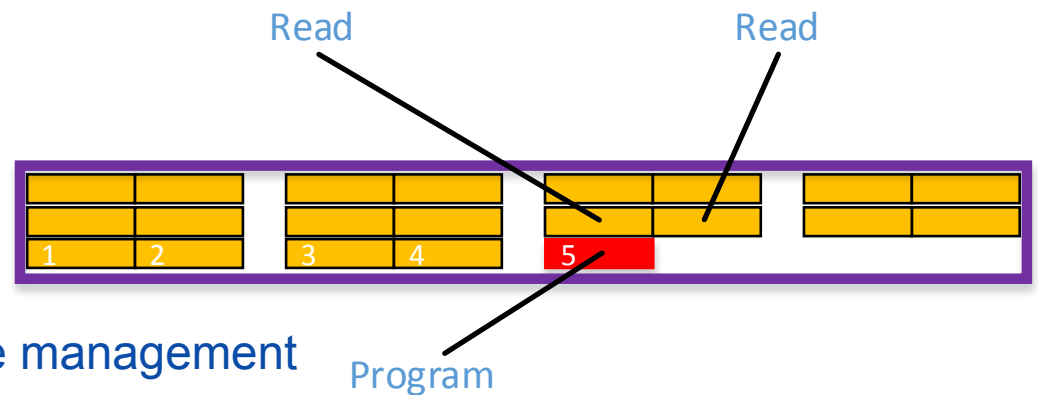
Burst Performance

- Host Burst Pattern
 - Logical Saturation
 - Size of the Cache
 - Rate of Recycling
 - Transition from Caching to direct Write
 - Consistent Throughput
- Endurance
 - Wear Management
 - SLC / MLC / TLC vs. QLC Wear Ratio



Write Throughput

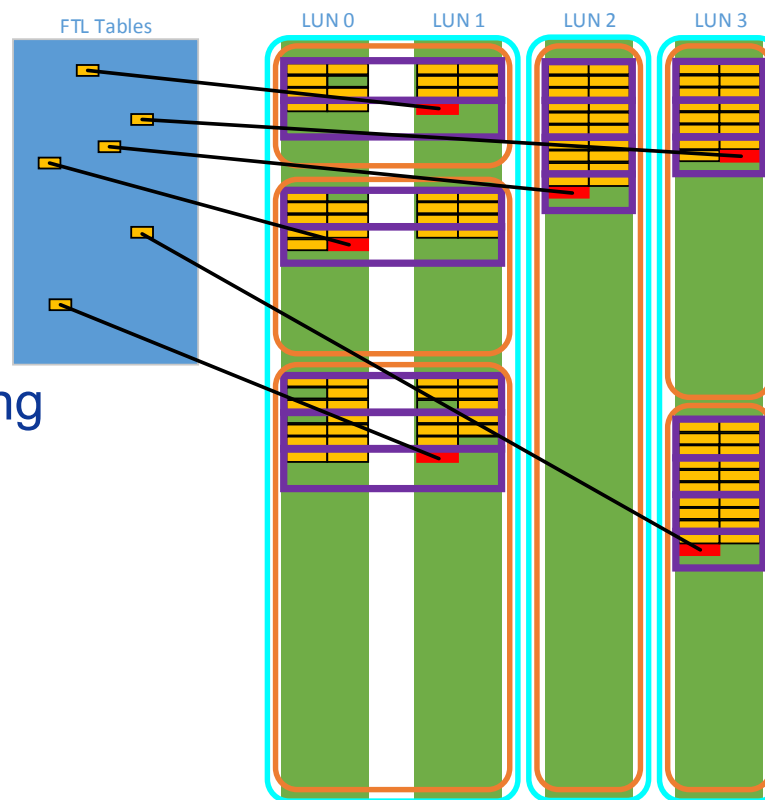
Active Management



- QLC Programming requires active management
- Tuning required for partially written Blocks
 - Erase History
 - Read History
 - Program Tuning
- Traditionally we leverage relationships between physical Blocks and Pages

Write Throughput

Subdivision of Media



- Physical Organization of Blocks becoming increasingly complex

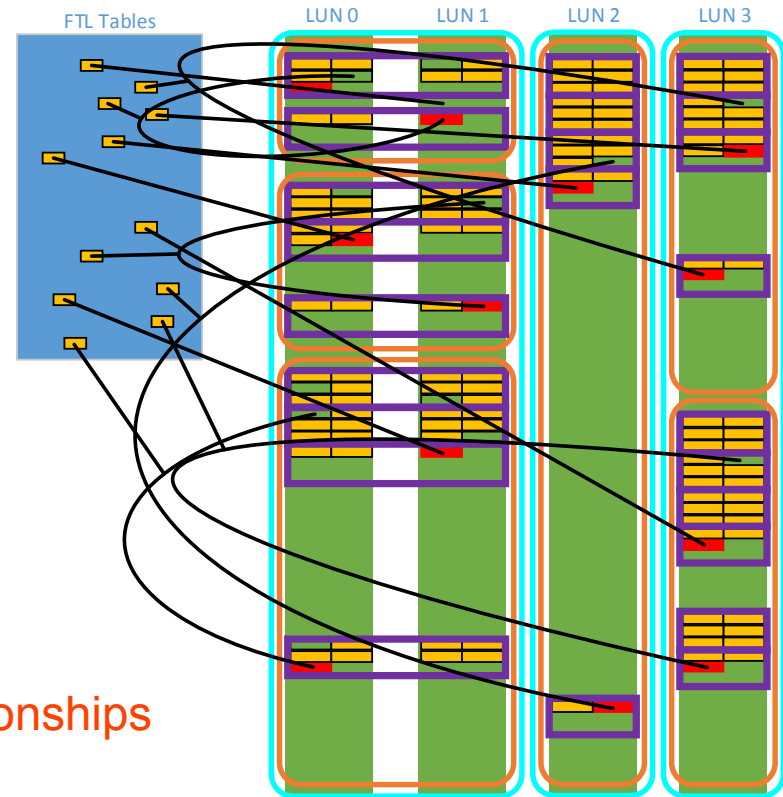
- Fewer LUNs
- Endurance Sets
- Namespaces
- Streams



Write Throughput

- Don't forget about Recycling
 - Doubles the number of active Blocks

As Physical Data Layout complexity increases,
We can no longer leverage Block to Block relationships






Read Latency

Minimal and Consistent

- Die Characteristics
- Block History
 - Erase Count, and Time
- Page History
 - Read Frequency
 - Adjacent Page operations
 - When and at what temperature data was written
- Significant Tuning Data Available
 - Pre and Post Computation / Organization
 - At a fine granularity
 - Accessed and updated Every IOP

Manage the Trigger Rate

Continuous Read Tuning






Flash Memory Summit

Summary

- Challenges and Efforts
 - Significant Logical and Physical data tracking
 - Continuous Tuning and Adjustment
 - Directed and Aligned NAND characterization
- Endurance
- Write Throughput
- Read Latency / Consistency



QLC offers many options over TLC, but is exponentially more complex to manage.



Flash Memory Summit

Thank You!