



Flash Memory Summit

DCTR-101-1: Enterprise QLC, Optimized

Tim Fisher

IBM STSM, FlashCore Module Chief Architect, Master Inventor

fisher@us.ibm.com

Enterprise QLC, Optimized

- Performance
- Reliability
- Cost



FlashCore Module

FCM3 Offering

2.5" Dual ported
NVMe SSD
U.2 Form Factor

Encryption with
FIPS 140-3 L2
Certification in
progress

Used exclusively
In IBM Storage
Appliances



Market Loading Density



World Class Characterization

IBM Research, FlashSystem, and Memory Development all work together to characterize and validate FCM memory technology

"Not all flash blocks are created equal"

Characterization helps to determine both differences and similarities in reliability and performance characteristics at both a block and page level

Exhaustive Characterization

Flash devices are continuously tested and analyzed to understand changes in *RBER due to PE cycles, retention, read disturb and much more...

*RBER – Raw Bit Error Rate


Secret Sauce

Algorithms developed due to characterization are then applied to the characterization data to verify we can meet and even exceed enterprise reliability and performance.




World Class Characterization...

- Thorough characterization is key to understanding what needs to be done to make QLC ready for the Enterprise
- Development and research work together to improve the nominal behavior
- When the Secret Sauce is applied, we then know if we are Enterprise ready...



Hardware Compression

- Inline, At Speed, Hardware Compression
 - Shipping since 2017 in FlashSystems products
- Data is decompressed along side compression and checked bit by bit. Never any corruption!
- Data compression/decompression algorithm is a Modified Dynamic GZIP
 - Performs well across a wide range of datasets and applications
- Takes advantage of already existing LSA mapping
- Decompression done in line with hardly any extra latency
- Data protection (ECC) is implemented on top of compressed data.
 - Allows garbage collection and other background data transactions to operate on compressed data
- Compression and Decompression completely transparent above the Flash module except for management of space




Compression increases performance, reliability, sustainability, and lowers cost



Hardware Compression...



- Compression Ratio
 - FCM provides LPT addressing for 3x the useable capacity (more for small)
 - However, FCM is capable of achieving much higher compression ratios, depending on the data
 - Data that is highly compressible will take up less physical space, which will increase performance and *further* extend the reliability of the FCM
- FTL Complexities with Compression
 - ECC Codewords (CW) are a fixed size and the same number of CW(s) will be written to each flash page
 - There is no alignment between the start or end of a compressed logical page and the start or end of a physical page
 - Many logical pages can reside in one CW or one logical page can reside in multiple CWs
 - Proprietary LPT mapping scheme allows FCM to efficiently pack and map logical pages without large metadata structures





Optimized FTL and Datapath

- Hard decision encoder
 - Read data from flash **ONE** time
 - Powerful ECC increases reliability
- Smart data placement
 - Data is packed to provide consistent low read latency across a wide range of queue depths
 - Data placement also allows the controller to take advantage of performance enhancing features
 - Multi-plane reads and writes
 - Sub-page reads
 - Read heat segregation






Consistent low latency and performance with increased reliability



Optimized FTL and Datapath...

- Calibration
 - IBM proprietary algorithm to track voltage shift
 - Read data from flash **ONE** time
 - Calibration done in background
 - Pages and blocks with similar characteristics are grouped to reduce metadata and reduce background activity
- Health Binning
 - Not all blocks are created equal
 - Blocks are graded during calibration using RBER
 - PE cycles alone are not an accurate representation of health
 - Selective data storage

Consistent low latency and performance with increased reliability





Computational Storage

- Inline compression provides *supreme* system processing offload and system memory offload
- Cutting Edge AMD / Xilinx FPGAs with embedded, powerful multi-core ARM processors
- FTL processing and metadata contained completely inside FCM.
Zero system resources needed!
- Additional Real-time Processing Units (RPU) and programmable logic available for potential future upgrades.
- Initial focus on Ransomware Detection



FCM is positioned to lead and shape the evolution of computational storage in All Flash Arrays and **beyond**...



IBM FlashSystem Products 2022

Entry Enterprise

FlashSystem 5200



Midrange Enterprise

FlashSystem 7300



High-end Enterprise

FlashSystem 9500 & FS9500R



IBM Spectrum Virtualize

Storage function, scalability, interoperability, cloud integration and automation



IBM Storage Insights

Anomaly detection, fabric monitoring, full stack visibility, predictive support

When you interact with IBM, this serves as your authorization to Flash Memory Summit or its vendor to provide your contact information to IBM in order for IBM to follow up on your interaction.

IBM's use of your contact information is governed by the IBM Privacy Policy.





Questions

Please take the Session Survey Thank you!



Tuesday, August 2



Wednesday, August 3



Thursday, August 4