



An Insider's Look at Designing Flash Controllers

- **by Holly Frost**
- Founder, CEO, Architect
- Texas Memory Systems, Inc.
- August 9, 2011
- Flash Memory Summit 2011



The CEO wants to increase his bottom line with Flash storage.

- He (she) wants the switch to Flash storage to be Quick and Painless.
- He wants a higher performance Computer system in order to bill more for computer time.
- He wants a more reliable Computer System.
- He wants to save Space and Power.
- He wants Easier storage management.
- He wants Flash storage to be better than Disks



Agenda

- Flash Controller Overview
- Flash Controller Design Decisions
- Flash Storage Performance
- Flash Storage Reliability
- Flash Controller IP Concerns
- Does Enterprise Still Need SLC?
- **Concept of a Flash Storage System**



Flash Controller Overview

- Flash Controllers determine features & performance – the controller is the “**Secret Sauce**”.
- Flash storage manufacturers are differentiated by their Feature Set in their Flash Controllers.
- Based on the Desired Market, designing a Flash Controller will involve dozens of decision points.

Key Controller Design Decisions

- Flash Storage Form Factor?
- Flash Capacity?
- Performance: Bandwidth, IOPS, Latency?
- Application: Read-Intensive, Write-Intensive?
- Buy or Design Controller?
- Embedded Processor?
- ASIC versus FPGA?
- Which Type of Flash to Use? (MLC, SLC)
- Flash Manufacturer? (T, S, M,,...)



Flash Performance

- Which Form Factor?
- How many Flash chips and Channels?
- What is Maximum Board Connection Speed?
- How does Flash Controller do heavy Writes?
- One or more Flash Controllers?
- On-Board CPU engine and Look-Up-Table?
- ASIC or FPGA Flash Controller?
- Your IP Controller Design or Licensed Design?
- What Is Board's Power limit?

Flash Reliability

- Flash chip manufacturers only guarantee that **98%** of the chip works.
- Flash chips must continuously be managed for new bad blocks and wear leveling.
- While the Flash data sheets mention bit errors (ECC), they do not mention other catastrophic (die or plane) failures.
- **Every new Flash generation must be fully characterized by the Flash board Designer.**



TMS Extra Reliability Features

- RAID across Flash chips.
- Perform Extra ECC for the Flash chip.
- TMS's Variable Stripe RAID to bypass bad Flash Planes or Dies.
- One month TMS Factory burn-in prior to shipping to reduce infant mortality Failures.
- Long term Flash device analysis to determine aging characteristics.
- Test and Select the Best Flash manufacturer.



Flash IP Concerns

- Is the Product easy to Reverse Engineer?
- Are you Protecting the “**Secret Sauce**?”
- Which new Ideas should be Patented?
- Which new Ideas should be Hidden?
- Which new Ideas can be used in Marketing?
- Who are you delivering the initial Product to?
- Do you encrypt your Programmable IP?
- Do you trust your Supplier and ASIC foundry?



Does Enterprise Still Need SLC?

- The answer depends on who is talking.
- SLC and MLC Flash have their advantages:
- **MLC is half the cost of SLC.**
- **SLC has 20x the Program/Erase Cycles.**
- The Enterprise MLC Market will still expand faster than SLC as MLC enters previous **SATA** space.
- **SLC still has an important place in Enterprise.**



SLC Enterprise Applications?

- SLC will be used for Write-Intensive Applications where Performance & Reliability are more important than cost.
- For the top storage Tier, SLC is the choice.
- For caching applications, its heavy write activity mandates the use of SLC.
- **SLC Flash will be used like SCSI drives, MLC Flash will be used like SATA drives.**

Concept of a Flash Storage

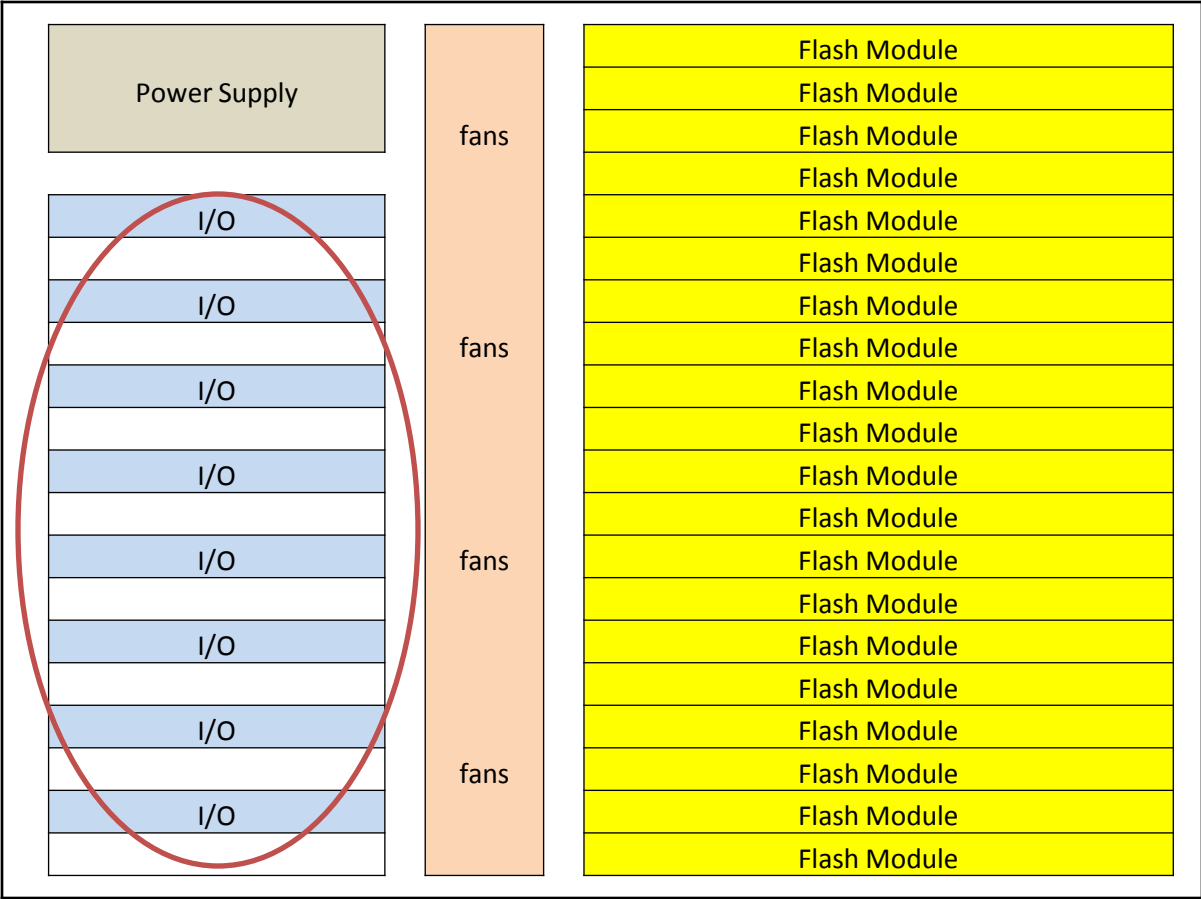
- **A Future Ideal Flash storage System:**
- High Capacity, Small Size
- High-Performance for Most Applications
- Rich Manage and Configuration Features
- Higher Reliability than Servers
- Auto Configuration Around Failed Storage
- Low Maintenance and Easy to Repair
- Low Power

Concept of a Flash Storage

- Form Factor = **4U Chassis Appliance**
- Capacity, Devices = **40-TB, SLC, 3200 devices**
- Controller = **TMS Designed FPGA chip set**
- System Bandwidth = **24 to 48 GB/s**
- System IOPS = **6m to 12m IOPS (4K)**
- Sustained Wt. Bandwidth = **8 to 16 GB/s**
- RAID = **RAID across Devices with VSR**
- ECC strength = **Factory Spec + extra bits**
- Production Burn-In = **30 days, minimum**




Concept Flash Storage






Concept Flash Storage





Concept Flash Storage





The Flash Developer Must:

- **Design Flash Storage that exceeds the Data Center Manager's Expectations.**
- Deliver the Flash Performance as Promised.
- Get better Sustained Flash Write Bandwidth.
- Execute Storage Management Internally.
- Increase Flash Density, Lower Flash Power.
- Make Flash storage more Reliable than Disk?
- **Keep your Customer's CEO Happy?**