



Intelligent Controllers for Best Performance

Bill Radke

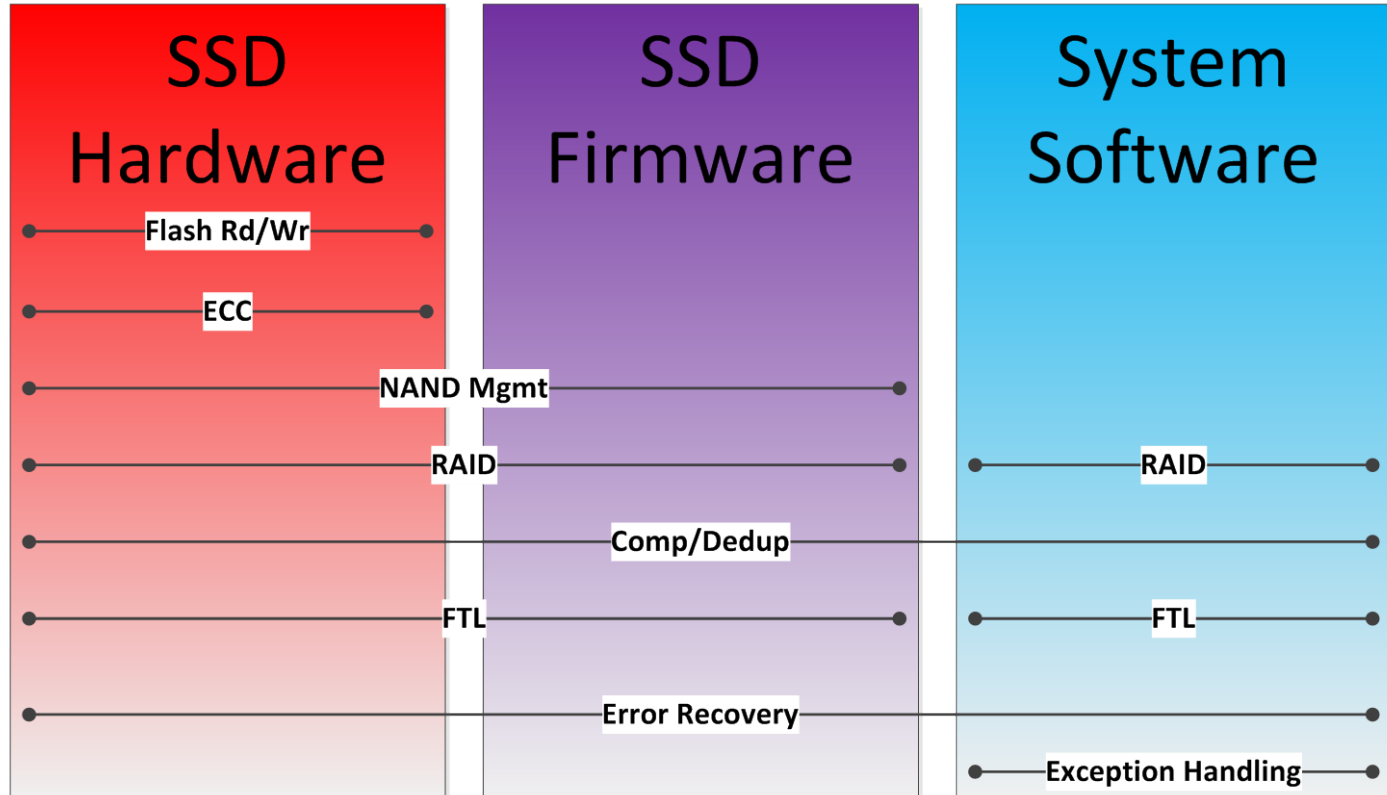
Director of Architecture

Skyera, Inc

Intelligent Controllers

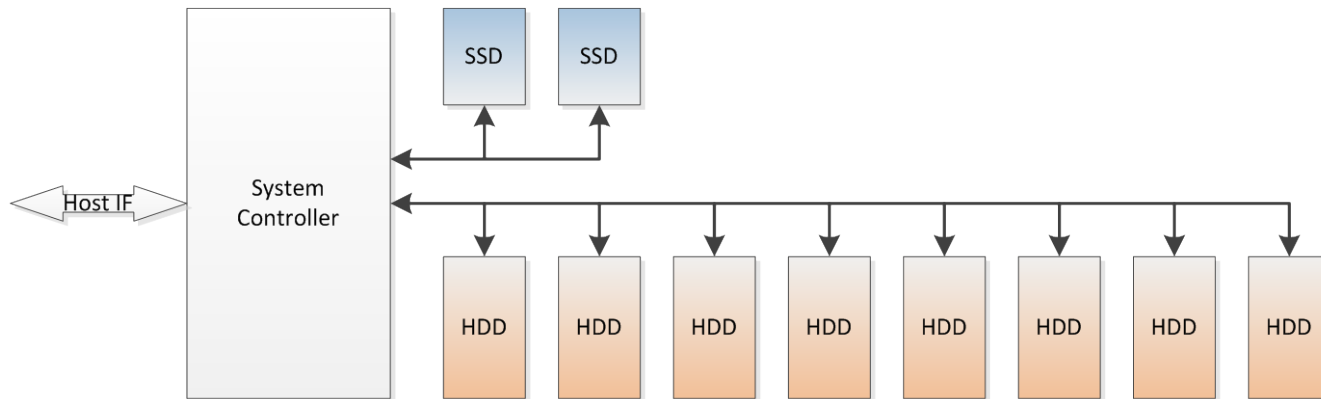
- Division of labour among software, firmware, and hardware
- System-based design, rather than SSD-based
- Key factor in optimizing system performance

Traditional Division of Labour




SSD-Based Design

- SSD's were designed with HDD interfaces, to fit into HDD systems
- Abstract NAND behaviour away
 - Main controller does not know what an SSD is doing



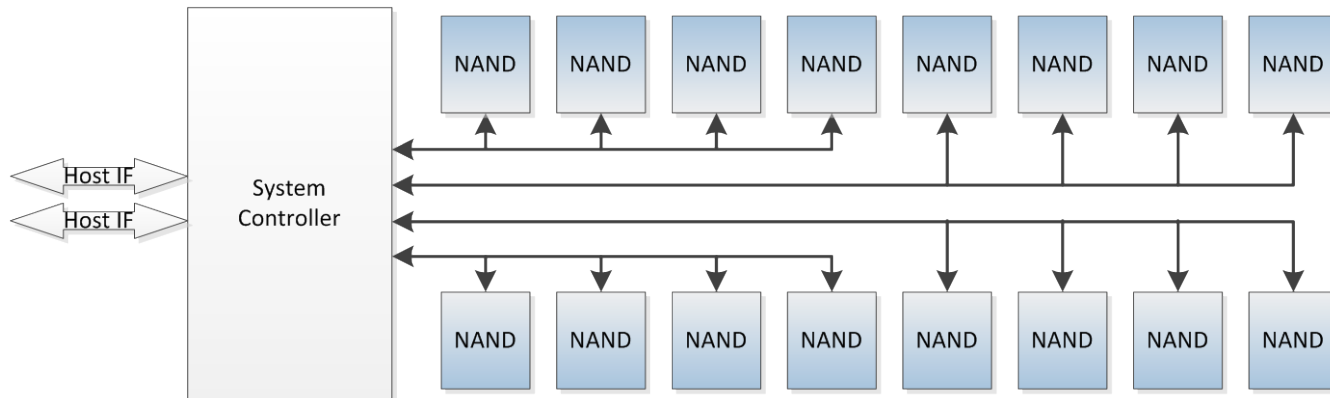
System-Based Design

- All-Flash Arrays can be structurally different
 - Optimize around NAND
- All parts of the design are NAND aware

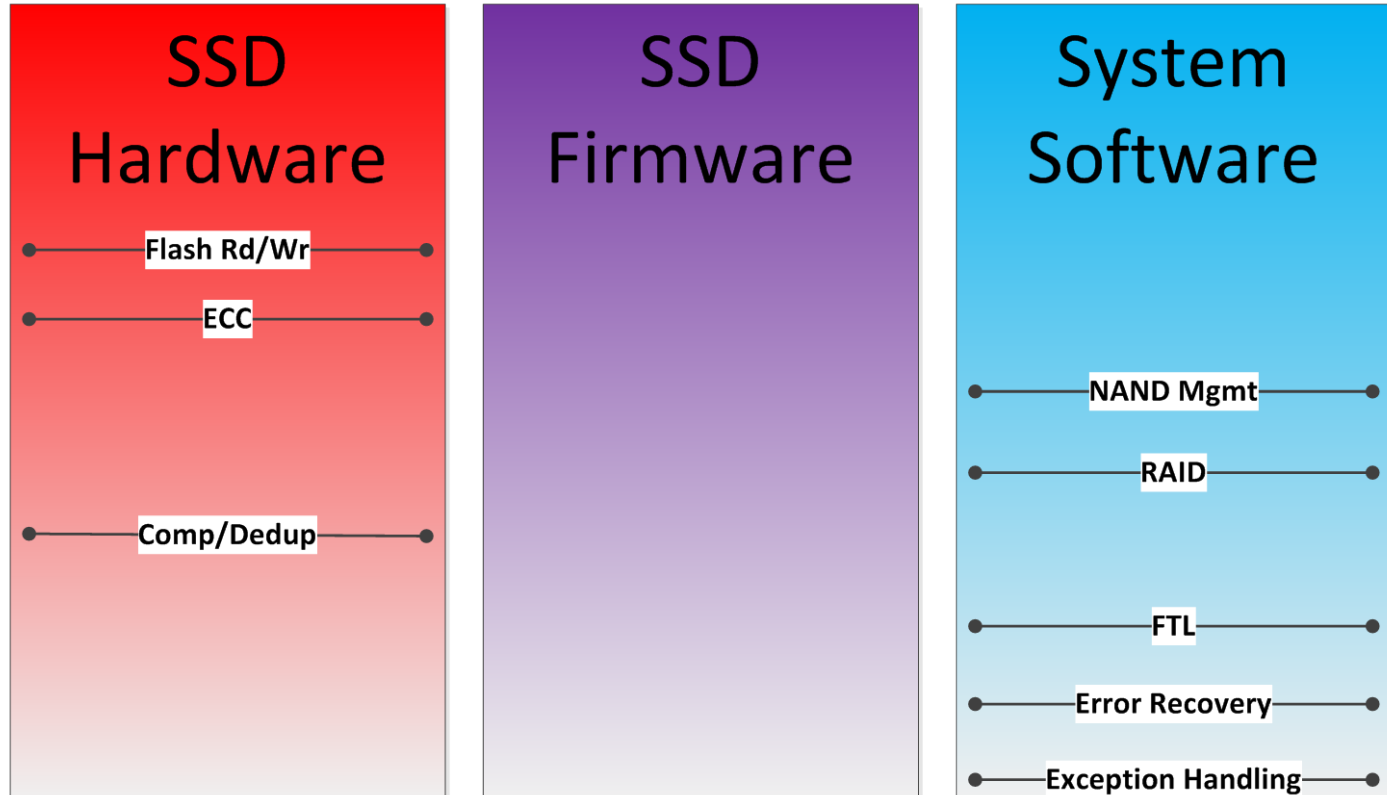


Intelligent Controller

- Control shared between NAND blades and the system
- FPGA-based blades perform hardware tasks
- Intelligence from the System Controller



Alternative Division of Labour



- Flash Controller handles hardware functions
 - DMA, read, write/program, compression, decomp, dedup
 - Lots of different command queues
- System Controller handles higher-level functions
 - FTL, NAND Management, Error Recovery

Problems with SSD-Based Design

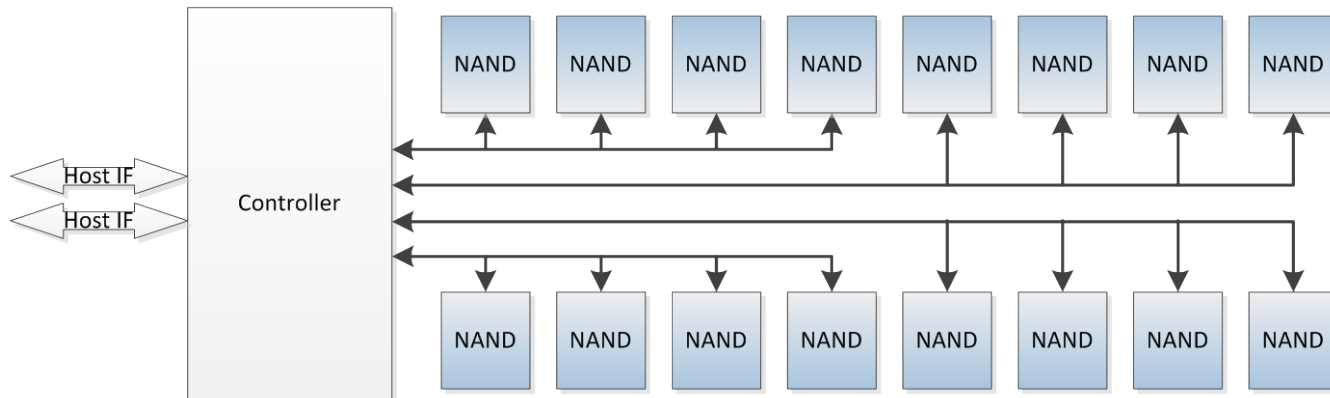
- Queue depth
 - High throughput means high latency
- RAID Structure
 - Two levels of RAID
- Page recovery
 - Variable response time
- Garbage Collection
 - Variable response time

Problems with SSD-Based Design

- Cost
 - NAND: \$0.38/GB
 - SSD: \$0.89/GB
- Density
 - Disk-based form factors
- Product Latency
 - Up to a year delay between NAND qual and SSD qual

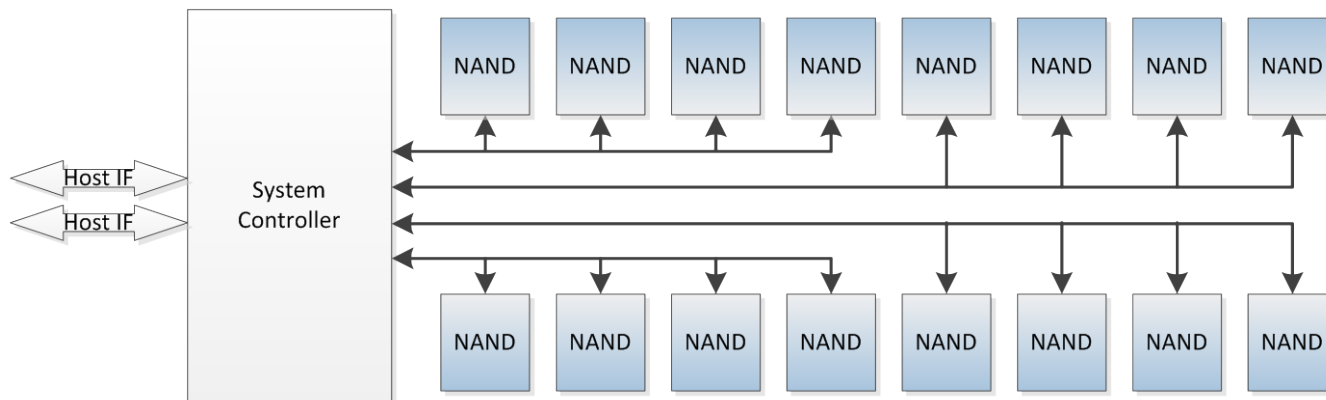
Intelligent Flash Controller

- Simple logic can fit into an FPGA
- Easy proto-typing, easy production
- Adapt to Most-Advanced NAND
- Adapt to new features





Intelligent System Control

- Software-based Flash Translation Layer
- Integrated error avoidance & recovery
- Centralized control complex
 - Multi-core CPU
 - Large DRAM arrays



System Performance


- Configurable
- Dense
- Fast
- Inexpensive



System Performance: Configurable

- Adjust to new features
 - Changed ECC algorithms
 - Updated page and codeword sizes
- Updated compression & dedup algos

- Support for Most-Advanced NAND
 - Up to 16nm 2D NAND
 - Expect to support 3D



System Performance: Dense


- Currently shipping 44 TB, 1U box
 - Highest density in industry
- Have shown 250 TB, 1U box



System Performance: Fast

- Blades give quick response to read, write, and program commands
- Accelerate in-line compress and dedup
- Spread traffic across all blades

- Current boxes can maintain 400k IOPS



System Performance: Inexpensive

- Buy cMLC NAND instead of eSSD
- Aggressively pursue the Most-Advanced NAND
- Minimize the system overhead
 - Components
 - Storage overhead
- < \$3.0 per GB
 - Formatted, no comp or dedup





Intelligent Controller: System-Based vs SSD-Based

- Improved system performance
 - Faster time to market
- SSD's will adjust
- Intelligent Controller should always be an option