Open Source Flash
The Next Frontier

Michael Cornwell

Lead Technologist - Flash Memory

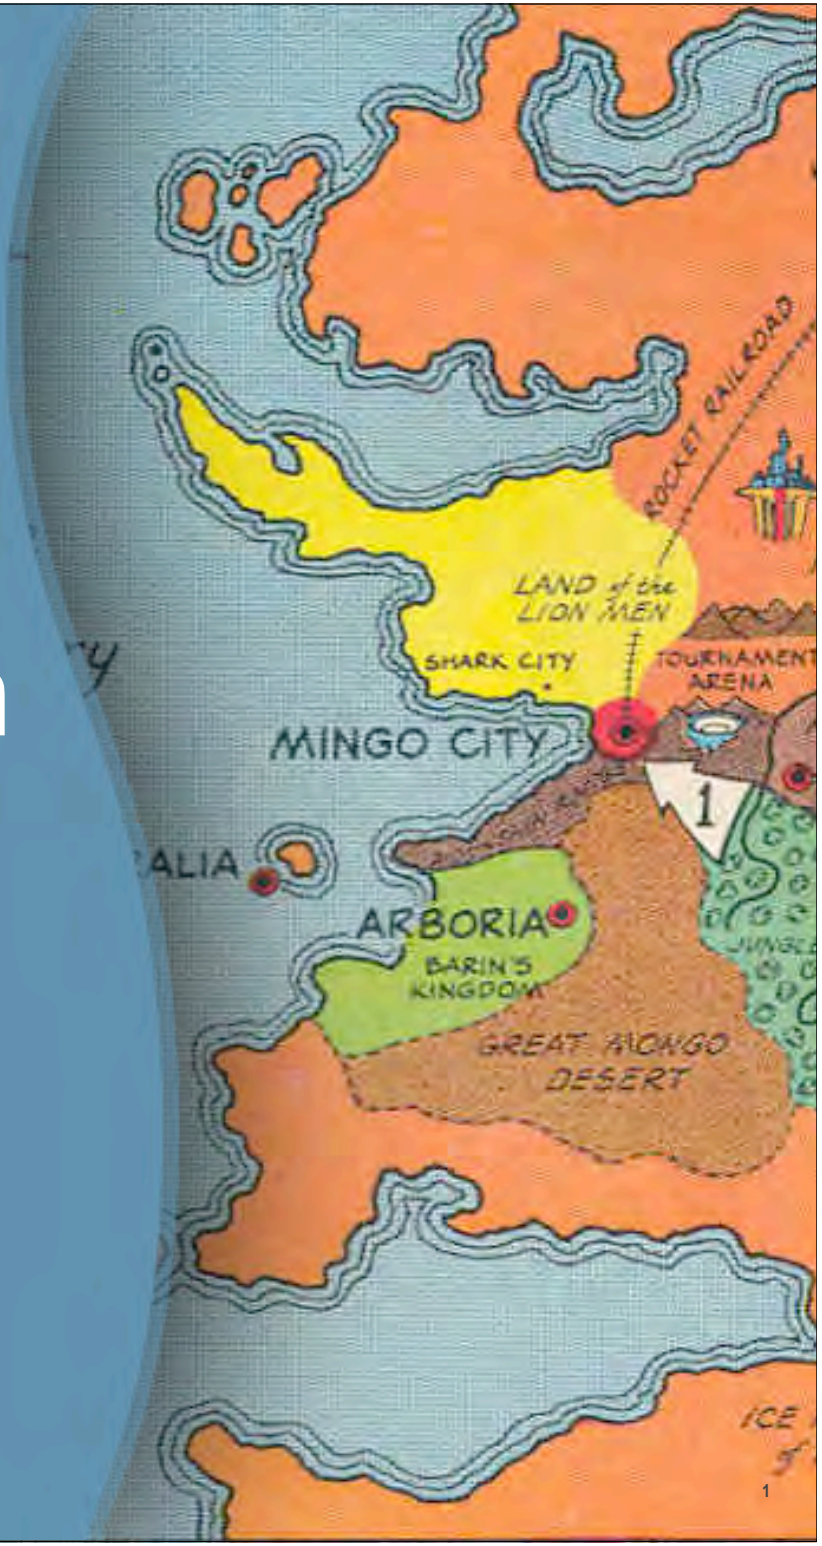

Flash for the Enterprise

A year after the great hype

- **Market Seen Too Small**
  - > Consumer is priority by NAND Suppliers
- **ASPs too High**
  - > SSDs cost more than systems they go into
- **Near Monopolistic**
  - > Economic issues
  - > Lack of execution from startups

Re-designing the Storage Element

- Most Efficient Footprint
  - > Every square millimeter matters
- Minimum flash required for performance
  - > Cost critical for adoption
- Leveraged Unused Chipset SATA Channels
  - > Free Host Controllers
- Aggregate at the channel level using software
- Focus on reliability and robustness of flash for 7x24 applications

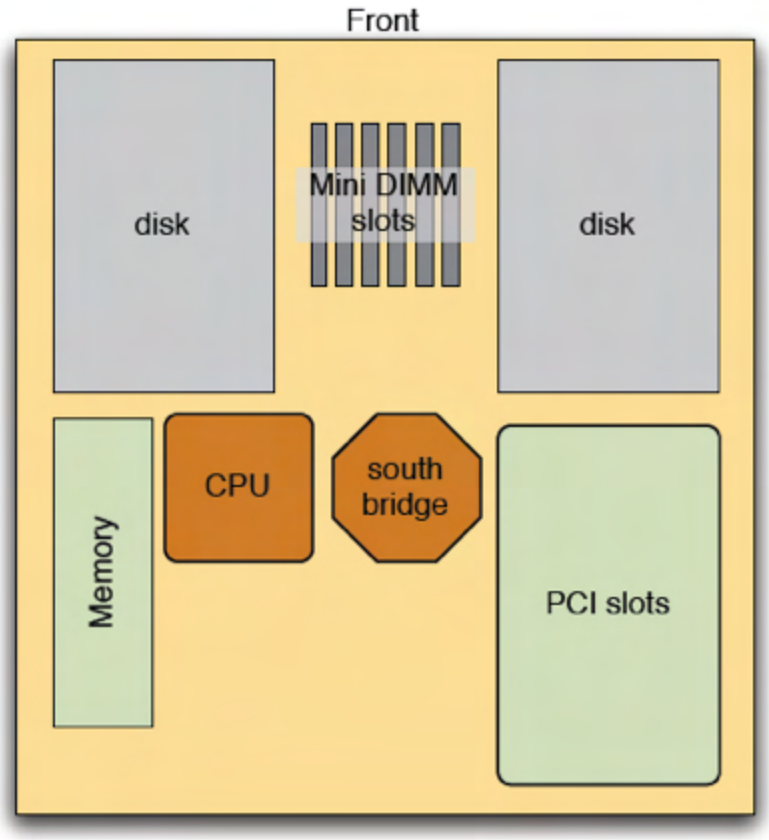

Original Open Flash Module Concept

Shrinking the IOP

Hard Disk Assembly

- 110MB/sec Read/Write
- 350 IOPS Read/Write
- 10 Watts
- 7502 mm^2

NAND Flash Die

- 40MB/sec Read/Write
- 7000 IOPs Read/146 IOPs Write
- 50 mWatts
- 165 mm^2

Open Flash Module

Next Generation in Server Storage

- First Server SSD designed for dense computing applications
- Minimum Size for Maximum Performance
- Leverages Existing Interfaces SATA/SAS
  - Reserved PCI-e in Pinout
- Released as an Open Standard to the Industry - JEDEC
- Supports optional backup power system

Evolution of Server-based Storage

2000

3.5” 15K HDD
16 Watts

2004

2.5” 10K HDD
8 Watts

2009

Open Flash Module
2 Watts

Modern Server Design

Pushing Thermal Limits to the Extreme

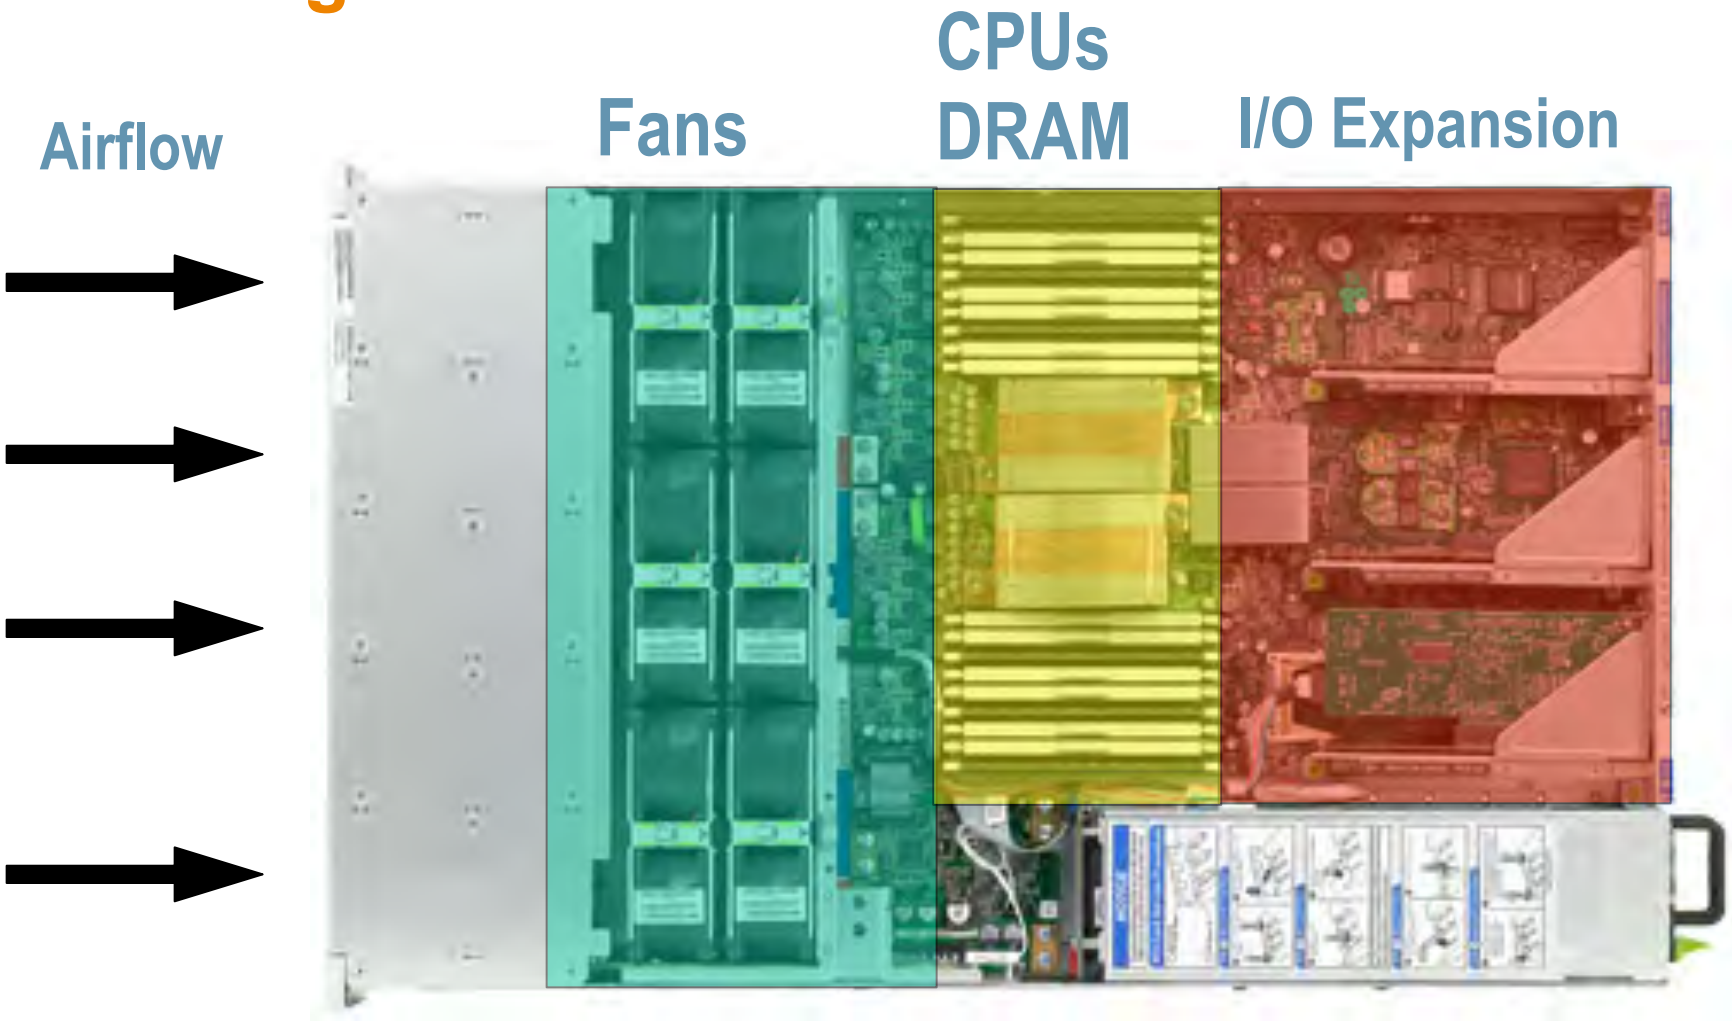

New Server Architecture

Flash Integration

- Flash inside for storage performance

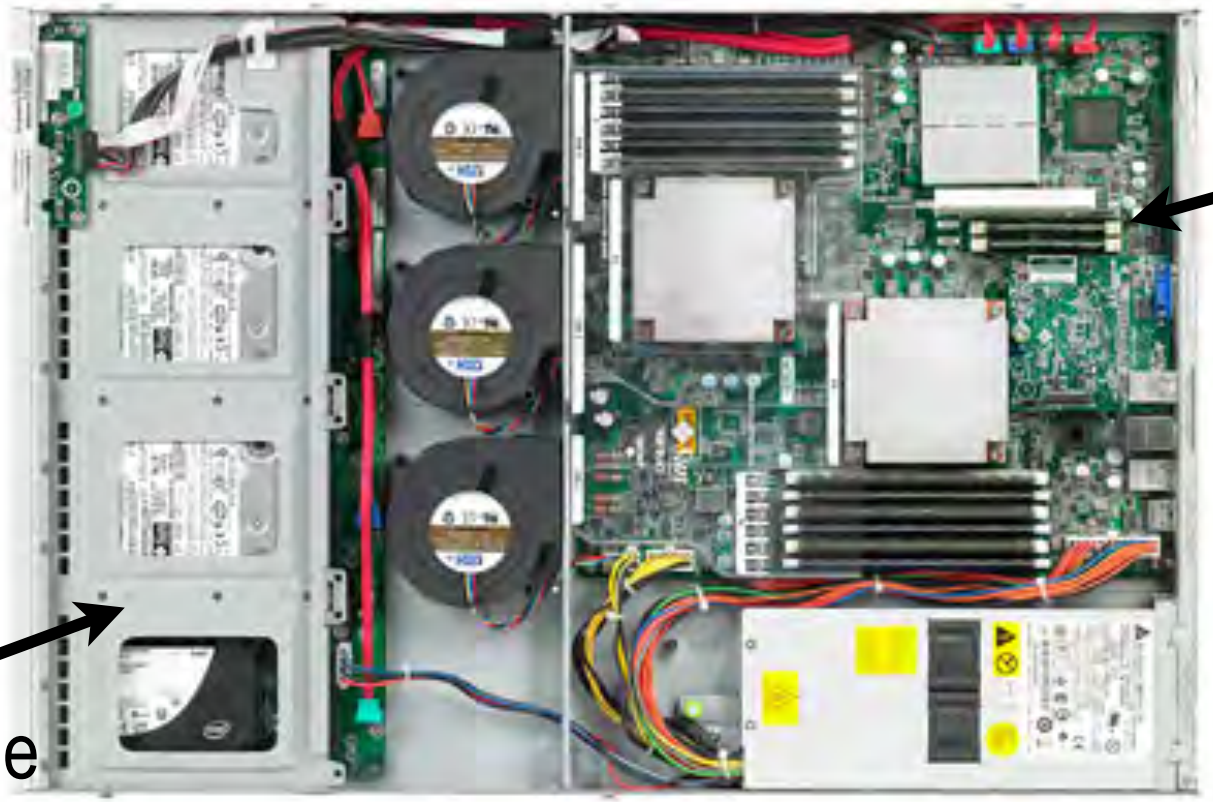

- HDD inside for storage capacity

Open Supply Chain Model

High Endurance
SAMSUNG
SERVER-GRADE SLC NAND

Unigen™
Solution For A Real Time World

MARVELL

- Source components independently
- No mark-up on components
- Tightly controlled manufacturing process
- Completely consumer defined

The $30 HDD

- Cheapest HDD to build
  - > 1 Head / 1 Platter
- Follows HDD Bit Density Growth
- Applications
  - > Game Consoles
  - > Netbooks/Low-end Laptops
  - > Cheap Desktop PC
- Price comparison today for $30
  - > HDD: 120GB ($0.25/GB)
  - > SSD: 16GB
- Technology Roadmap for 10-30x storage density

The Optical Drive

- Optical drives are included in systems today mostly to support software install and backup
  - > Tens of dollars in cost to system design
  - > Precious system design space
- Several Systems today don't include ODDs *(i.e netbooks)*

Software Distribution

- Distribute Software as a read-only device
  - > No Viruses
- O/S and Bios already support optical drives
  - > 100% emulation of CD/DVD drive
- Less packaging
- Perfect Application for 3/4bit MLC
- <$5 unit cost

*Image from: http://www.portablegadgets.net/index.php?id=116

The Lithography Dead March

Compromising the Value of NAND for Cost

- As NAND Manufacturers chase Cost per GB it is coming at a cost
  - > Endurance - 1/10th of 3 years ago
  - > Write Speed - 1/4 performance of 2 years ago
  - > Read Speed - 1/6 performance of 2 years ago
- NAND will have higher latency than HDD in 2 Generations
- Very few sub-50nm designs capable of supporting Enterprise applications
  - > Non-endurance failures biggest issue

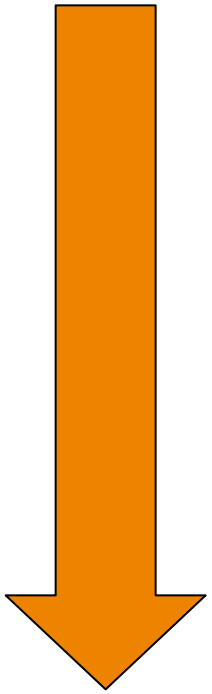

Non-Volatile DRAM
NAND backed DRAM Solutions

- Leverages Performance of DRAM with Non-volatility of NAND
- Attaches to existing high speed, low latency DRAM interfaces
- Applications Today
  - > Industrial / Military Applications
  - > RAID / Storage Controllers
- Requires software awareness of non-volatility
- Move towards high capacity PCI-e implementations in future

NV-DRAM Memory Module

NVMHCI - Enterprise Extensions

- Non-Volatile Memory Host Controller Interface
  - Fresh take on interface for SSDs and caches
  - Significantly lower latency than legacy disk interfaces
- Enterprise Features
  - Hot-plug Support
  - Multi-path Interface
    - Active/Active failover
  - Port Virtualization
  - NV-RAM capability
- True Replacement for SAS/Fiber Channel for Solid State

NVMHCI - Enterprise Extensions

Chassis Based Shared Memory

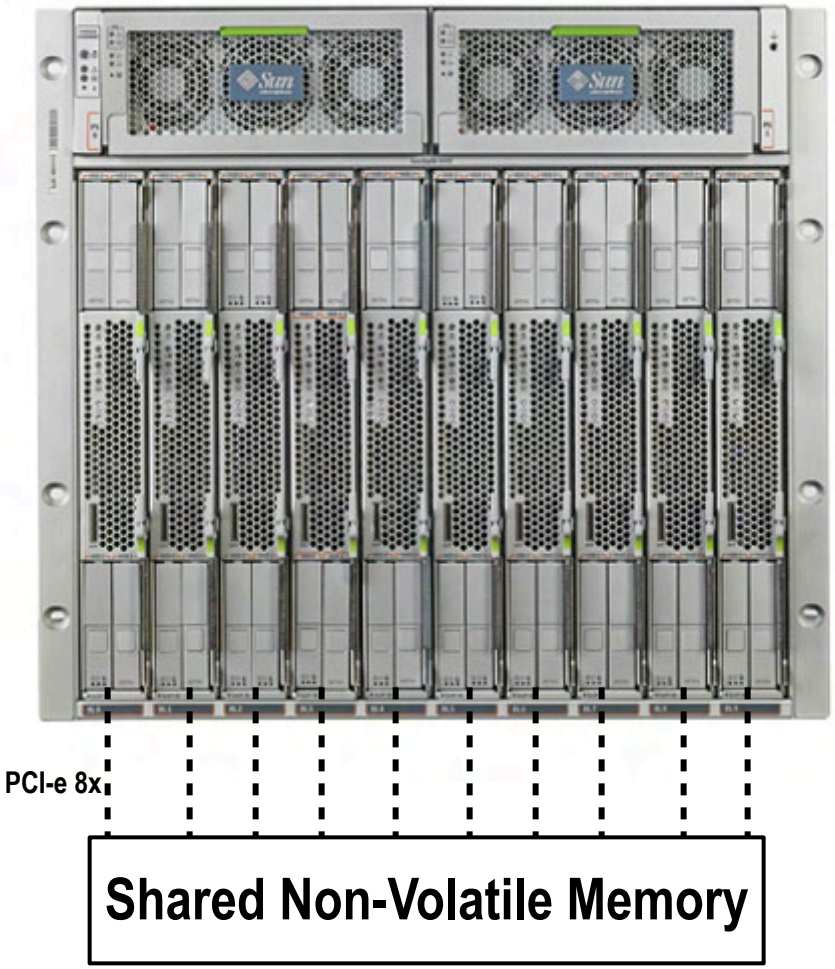

- PCI-e Interconnect
  - > Multi-path
  - > I/O Virtualization
- Tiered Memory Architecture
  - > 100's GB of NV-DRAM
  - > 10's TB of NAND
- Shared Programming Model exists Today
  - > memcached, OpenMP, VMCI
- Based on Open Standards

Thank you

Michael Cornwell
michael.cornwell@sun.com