All Flash Array Data Protection Schemes

Sean R. Barry

XtremIO System Engineer

All Flash Arrays – How Did We Get Here?

- 30+ years of HDD and distributed data and parity
- Built for spinning media
  - Data Protection built into the distribution schema

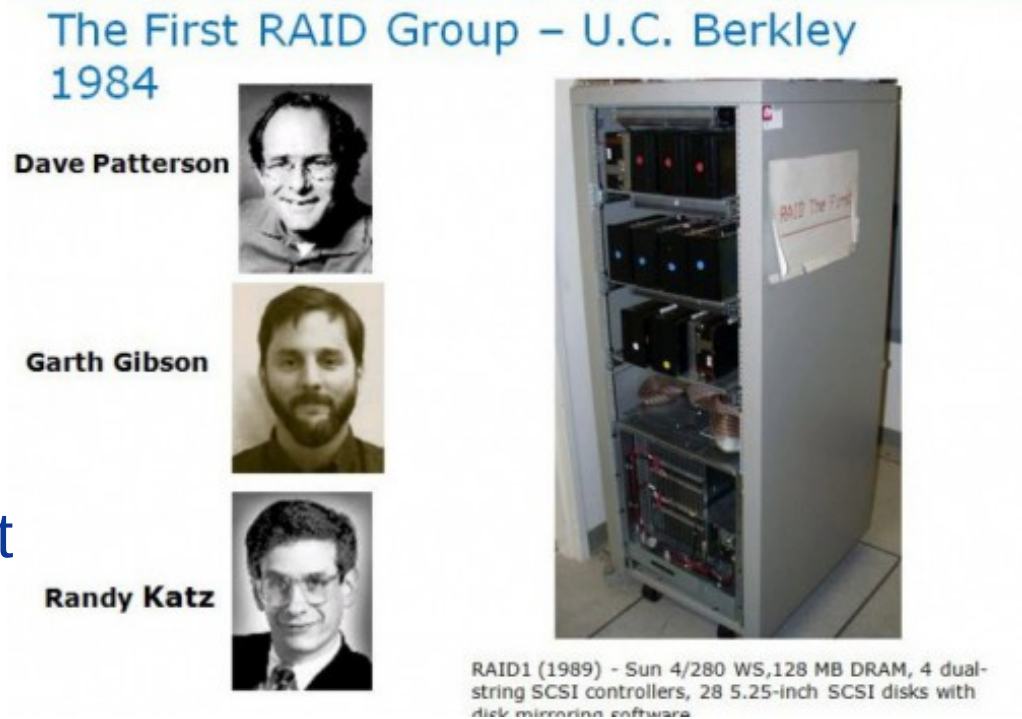

Source: gabrielchapman.com

All Flash Arrays – How Did We Get Here?

- Gen 1 - EPROM NAND 35+ Years old
- Gen 2 – Flash NAND 18+ Years of innovation

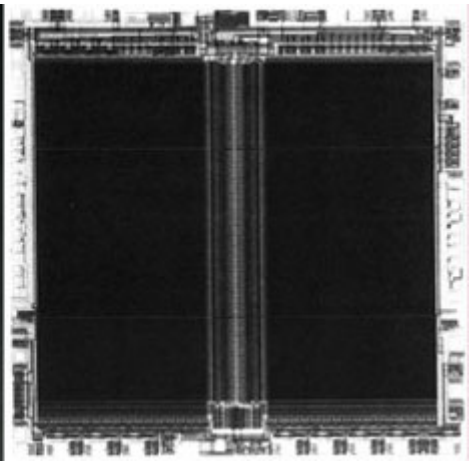

Source: spectrum.ieee.org

All Flash Arrays – How Did We Get Here?

- Gen 2 – Flash NAND 18+ Years of innovation:
  - Substrate
  - Gates
  - Bit Density
  - Manufacturing

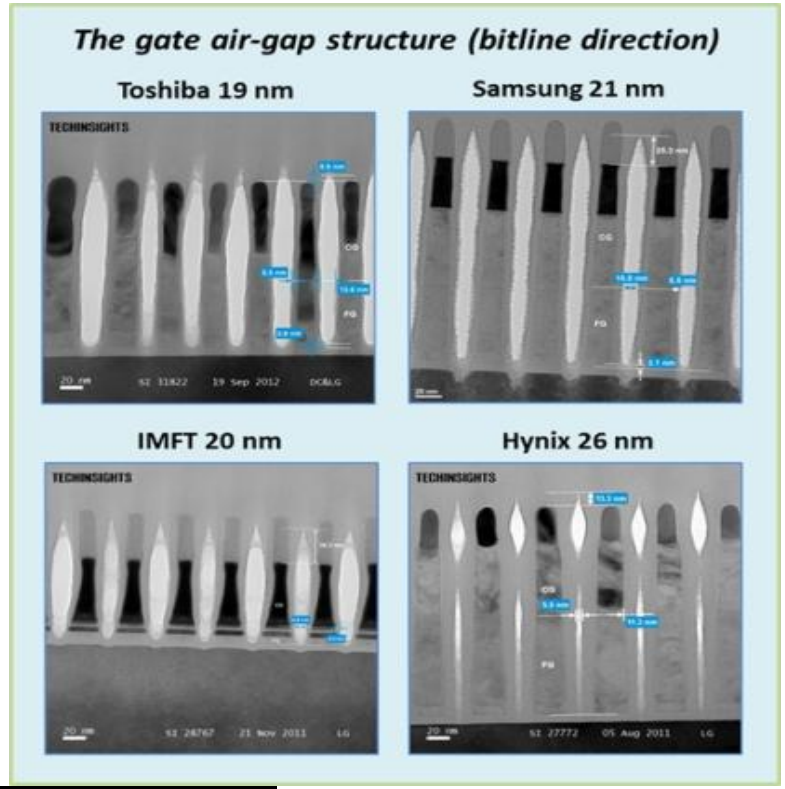

Source: www.eetimes.com

Hard Disk Drive Performance

Flash NAND Technology has an opportunity!

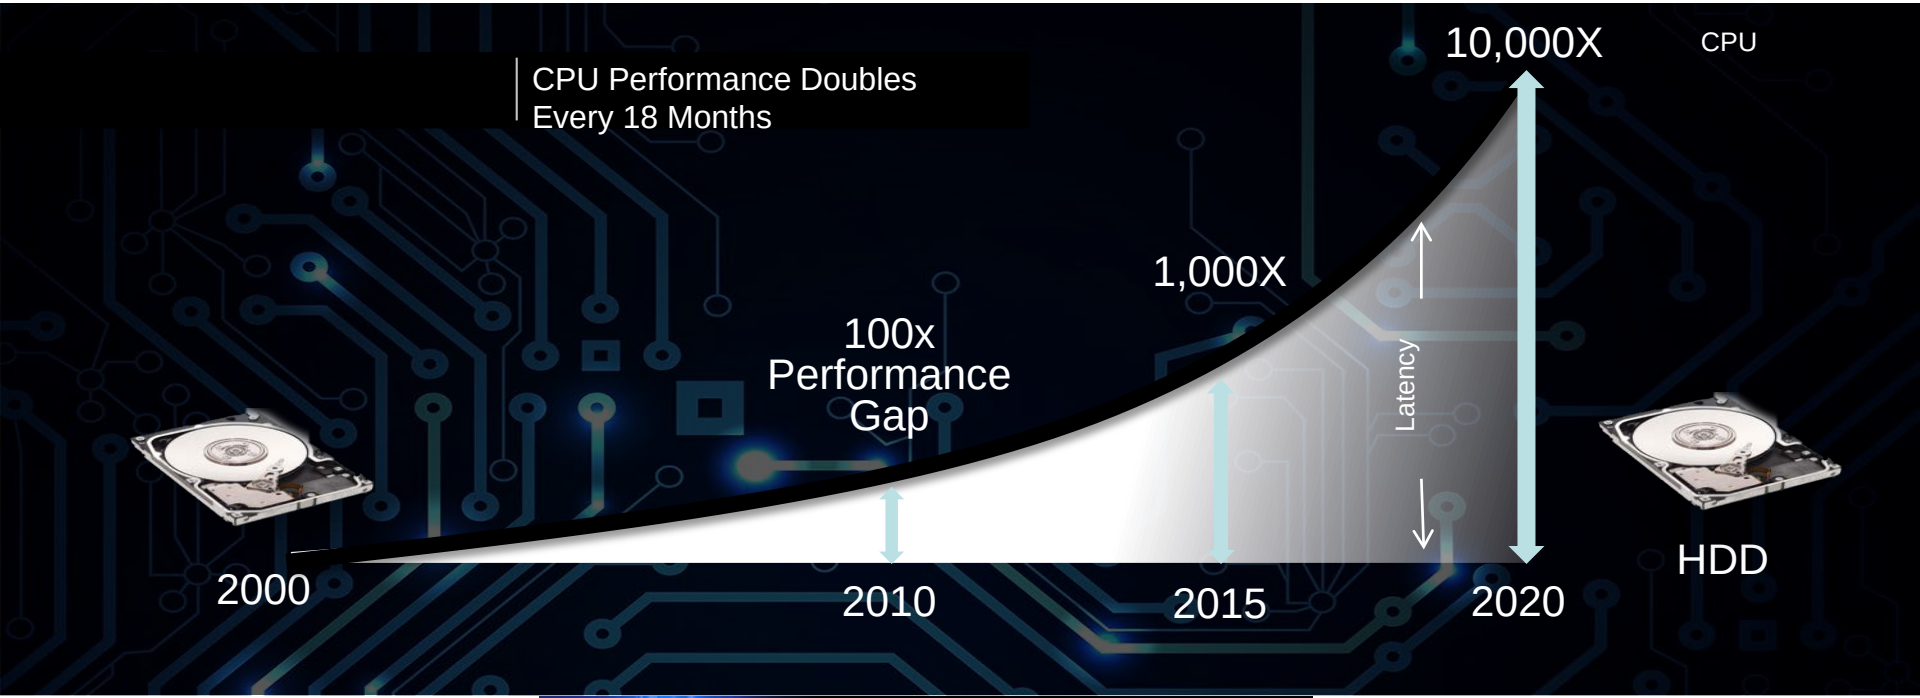

THE EVOLUTION OF “FLASH” ARRAYS

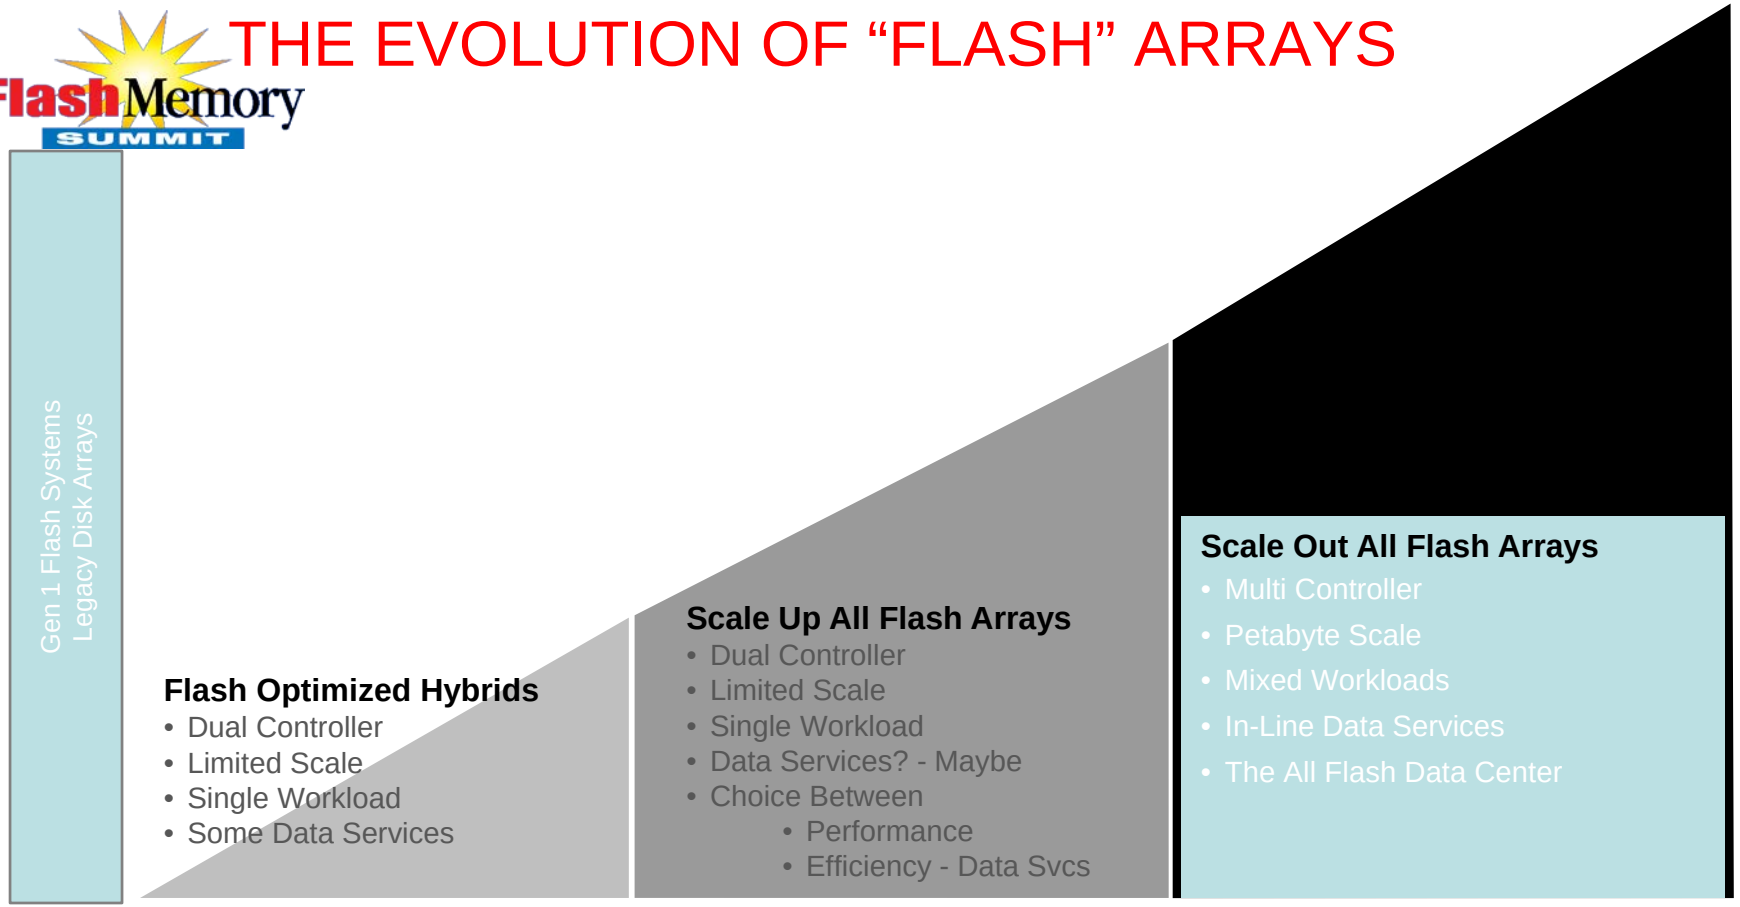

Gen 1 Gen 2 Gen 3 Gen 4

Source: Wikibon December 2014

ON LOAN FROM THE DISK ERA

Log Structuring	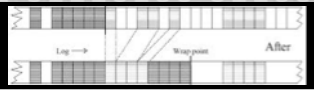	Garbage Collection Nightmare
Post-Processing		Erratic Performance
Legacy RAID		Write Amplification
Scale Up Architecture		Front End Bottleneck

Data Protection Methods on Flash

- Flash-optimized data protection without compromise

	XDP	RAID 1	RAID 5	RAID 6
Capacity Overhead	8%	50%	25%	20%
Write I/O Overhead (stripe update)	1.22	2 (64%)	2 (64%)	3 (146%)
Read I/O Overhead (stripe update)	1.22	0	2 (64%)	3 (146%)

SUPERIOR
USABLE CAPACITY

UP TO 4X
BETTER ENDURANCE

UP TO 4X
BETTER PERFORMANCE

Data Protection Innovation Shift - Hardware

- Lessons learned from enterprise storage field deployments
- Innovation at controller and HDD levels

- When RAID failed, manual data stripe rebuilding

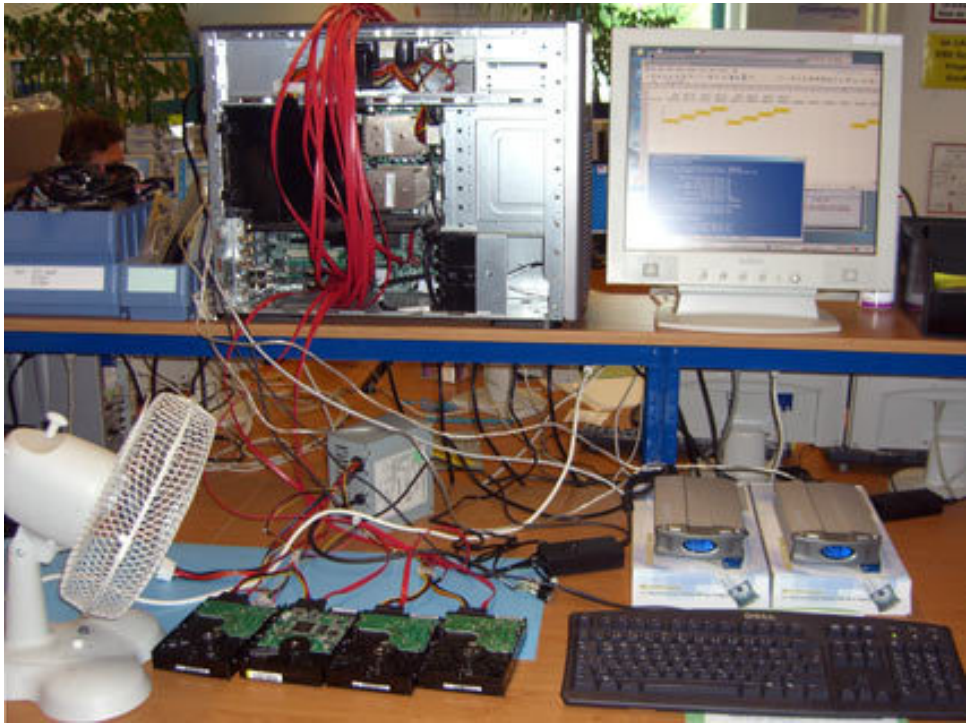

Source: www.tomshardware.com

Data Protection Innovation Shift - Hardware

- Hardware Architecture completely different because of Flash NAND Media
- Dual storage controllers
- Software driven data protection

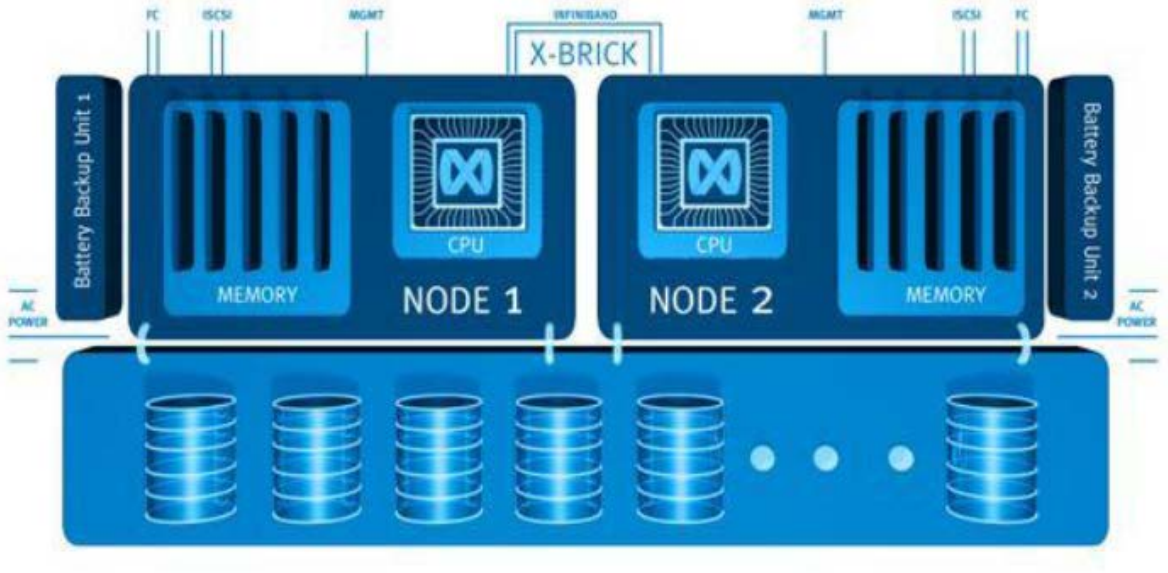

Flash Memory SUMMIT

Overview

Data Protection Innovation Shift – Metadata and Database Transaction-like I/O Operations

Emptiest Stripe Always Known

LBA offset	Fingerprint
100	963FE7B
104	CA38C90
108	134F871
112	0325F7A
116	963FE7B
120	F3AFBA3
124	AB45CB7
128	20147A8

Volume LBA to Fingerprint
(first stage metadata)

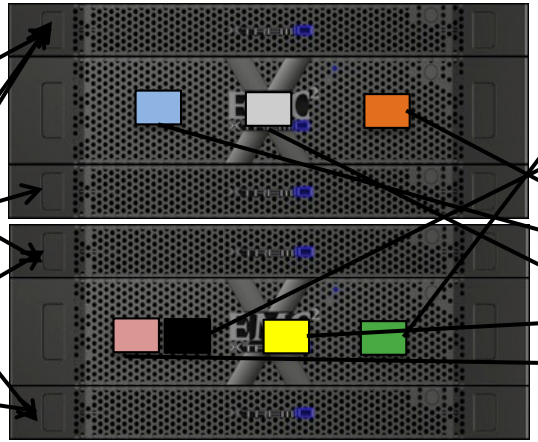

Fingerprint	Stripe #	SSD#	Ref. Count
20147A8	597863	1	1
AB45CB7	597863	12	1
F3AFBA3	621135	16	1
963FE7B	747322	19	2
0325F7A	143565	24	1
134F871	935843	10	1
CA38C90	342224	12	1

Fingerprint to Physical Location
(second stage metadata)

All Flash Array Take-Aways

- It's not about the Flash/SSDs—It's all about the architecture of NVM management
- All-Flash Arrays are more memory intensive to manage I/O and provide data services
- System design is key to All Flash Array longevity and datacenter utilization

Thank you!

Sean.Barry@emc.com