Education

NAND Flash Solid State Storage Reliability and Data Integrity -- an In-depth Look

jonathan thatcher, Fusion-io

SNIA Legal Notice

Education SNIA

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Abstract

- **NAND Flash Solid State Storage Reliability and Data Integrity**

  Abstract: "This tutorial provides an in-depth examination of NAND Flash as a storage medium and control techniques used to ensure enterprise-grade Solid State Storage (SSS) reliability and data integrity. Several aspects of the topic are discussed, including reliability characteristics unique to the medium (e.g. wear out, shelf life, infant mortality, failure modes), and implications for architecture, implementation, system integration, and deployment of NAND Flash-based SSS. These topics will be reviewed with a focus on device failure rates, data integrity, and data availability. Finally, recommendations are made regarding best practices, and information that users should obtain from potential vendors to properly assess NAND Flash SSS device reliability and data integrity."

Media Reliability / Availability

- The GOOD:
  - No moving parts
  - Catastrophic device failures are rare (post infant mortality)
- The BAD:
  - Relatively high bit error rate, increasing with wear
    - MLC wear rate (higher capacity density) worse than SLC
    - Higher density NAND Flash will increase bit error rate
  - Program and Read Disturbs
- The UGLY:
  - Partial Page Programming
  - Data retention is poor at high temperature
  - Infant mortality is high (large number of parts…)

Controller Reliability Management

- Wear leveling & Spare Capacity (e.g. Spare Blocks)
- Read & Program Disturb Controls
- Data & Index Protection
  - ECC Correction
  - Internal RAID
  - Data Integrity Field (DIF)
- Management

Poor Media + Great Controller → Great SSS Solution

Note: Multipage Programming Should Not Be Done

Quality Quiz

Education SNIA

Is the following example about

Performance

Or

Reliability?

Company: Wine.com

Education SNIA

- On-line internet Retail
- Representative Markets:
  - Transaction Processing
  - Data Mining
- Problem (4Q07)
  - Systems not capable → Contract out data mining
  - Performance hitting 100%→ Loss of revenue
  - Cost to meet growth: prohibitive

Wine.com

Before (3Q2008)

Internet

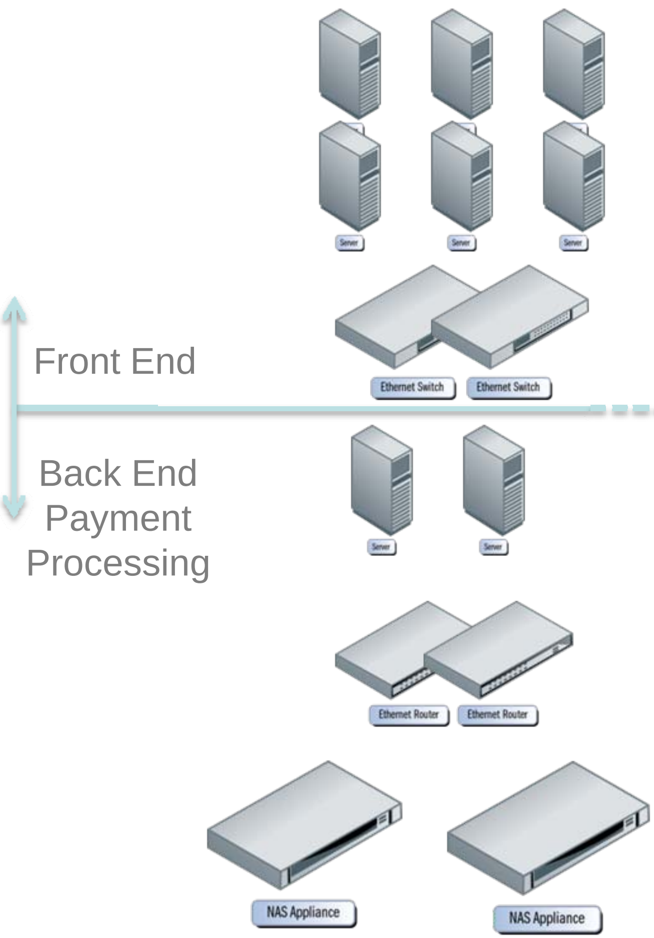

After (4Q2008)

Internet

Application level replication for redundancy

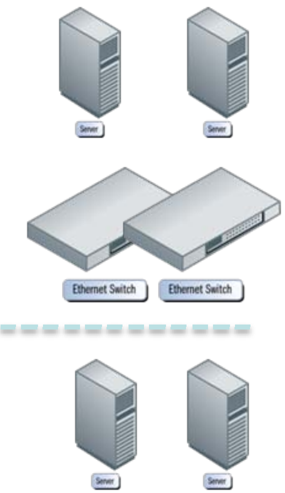

33% CapEx avoidance
50% OpEx reduction
33% Footprint reduction

12x improvement on write
- Latency down from 4 to <1ms

14x improvement on read
- Latency down from 12 to <1ms

“Enough capacity to cover 24 months of growth”

Geoffrey Smalling - CTO

RAS Improvements

Education SNIA

Change	Before	After	Reliability Gain
# Servers	8	4	Substantial
# Software Licenses	8	4	
# Ethernet Routers	2	0	Substantial
# NAS Appliances	2	0	Substantial
# HDDs	28	0	Substantial
# Ethernet Switches	2	2	
# SSS Devices	0	6	(Minimal)

And some think SSS is only about performance

Think Again!

Performance

↓

Reliability

Product Specifications

Education SNIA

Features directly affecting performance measurements

	SATA (A)	SATA (B)	PCI (C)
Capacity (GB)	32	32	160
Bus/Link	SATA-II (3 Gb/s)	SATA-II (3 Gb/s)	PCI-E X4 1.1
Memory Type	SLC	SLC	SLC
Adjustable Reserve Capacity	No	No	Yes
SSS Internal RAID -- Running during test	No N/A	No N/A	Yes Yes
K-IOPS (RMS)	8	27	88
K-IOPS (RMS) / WATT	3	?	7
Bandwidth (RMS, MB/s)	56	208	743
ECC correction	7 bits in 512B	?	11 bits in 240B

IF

product **performance** is **10X**,

can the **life** be **0.1X**?

Can **infant mortality** be **10X**?

Flash Memory Cell Wear-out

- With repeated P/E Cycles, electrons get “stuck” in the oxide, preventing a reliable “read” of a logical 0 and a logical 1.

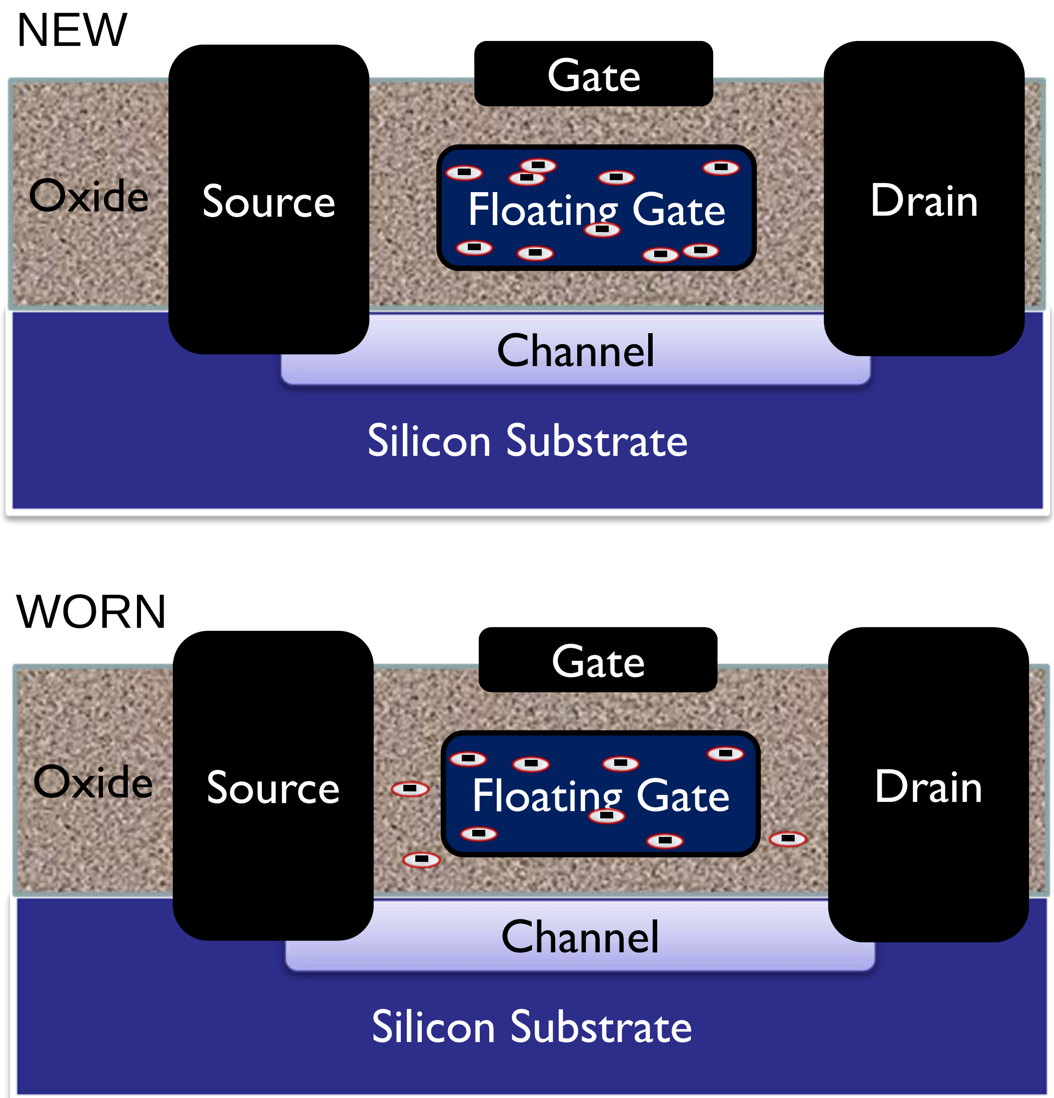

Memory Cell Wear Out

Higher BER

< Data Integrity

< Life

Raw Erase Error Rate (SLC, MLC)

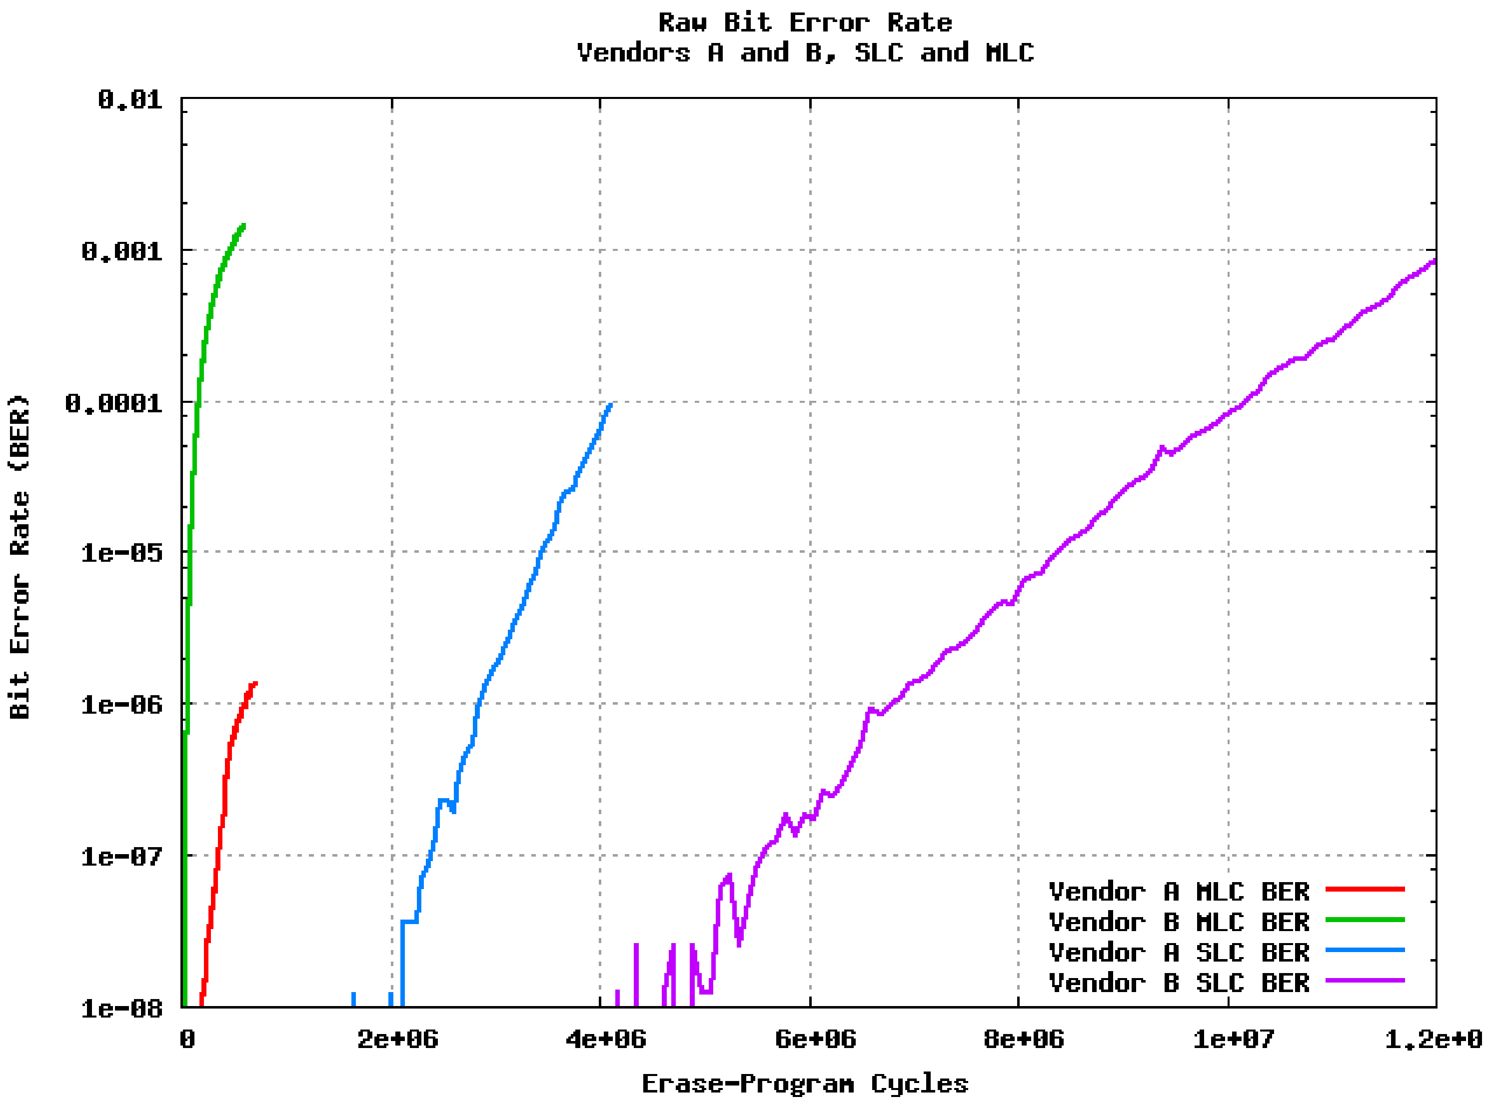

Read Disturb Raw Error Rate (MLC)

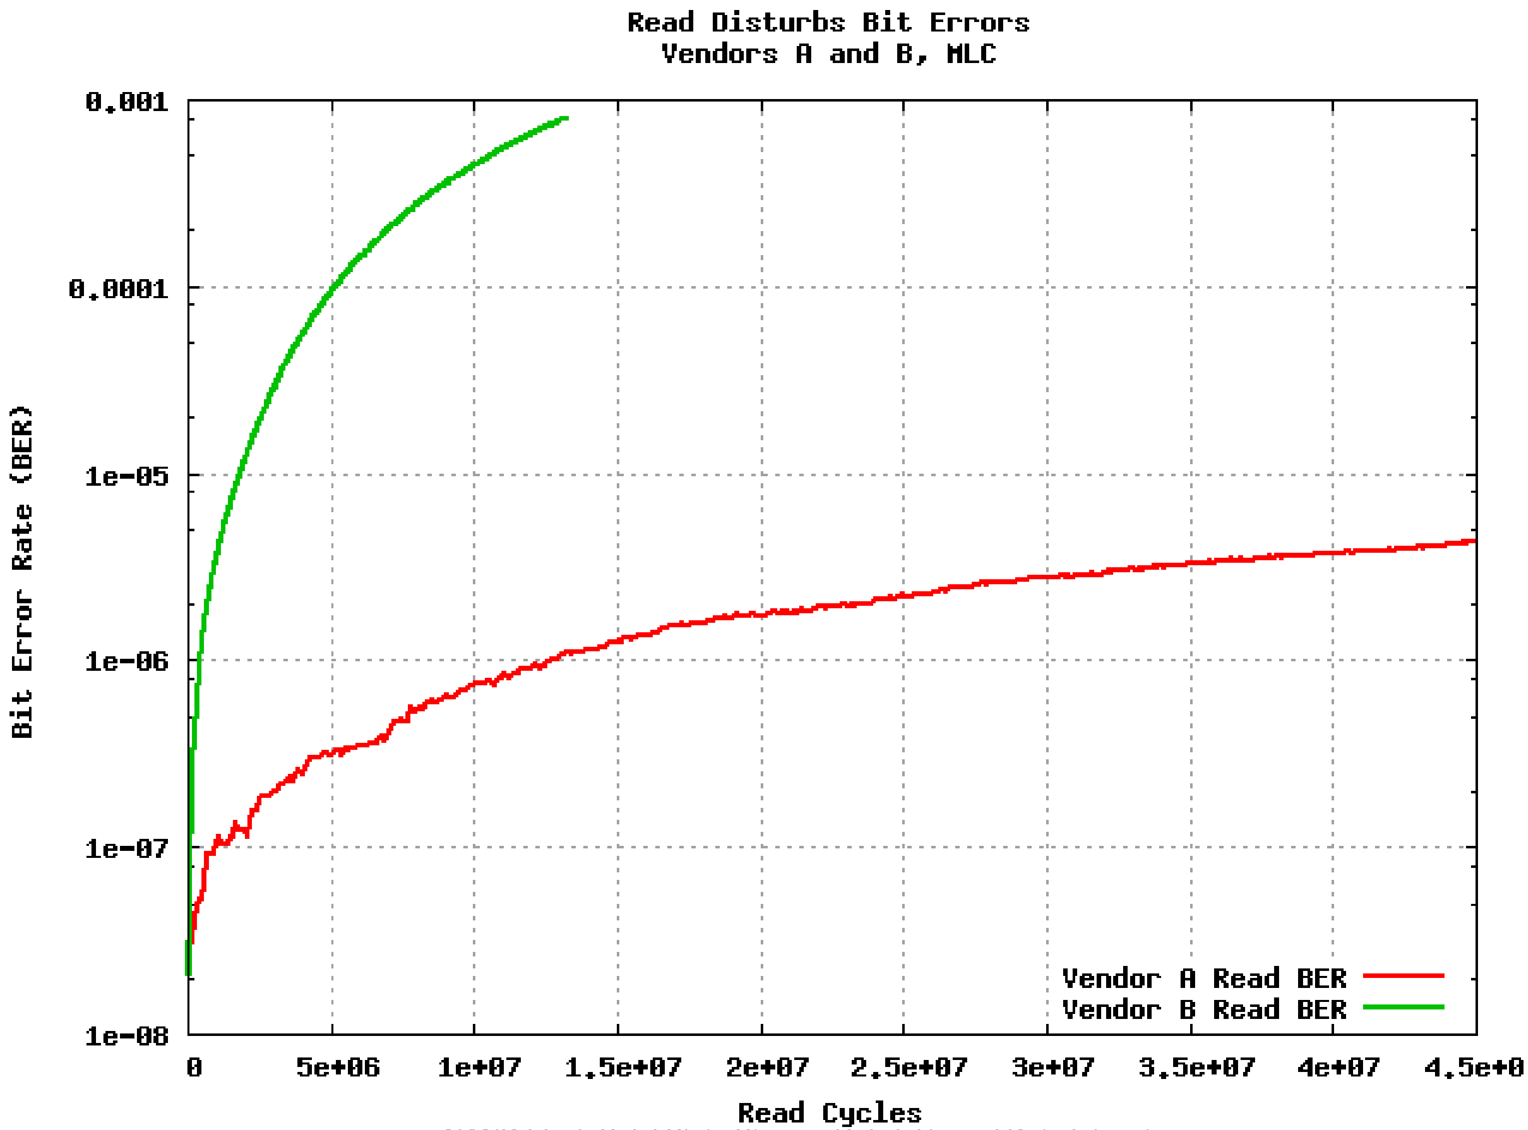

Retirement v Temperature (MLC)

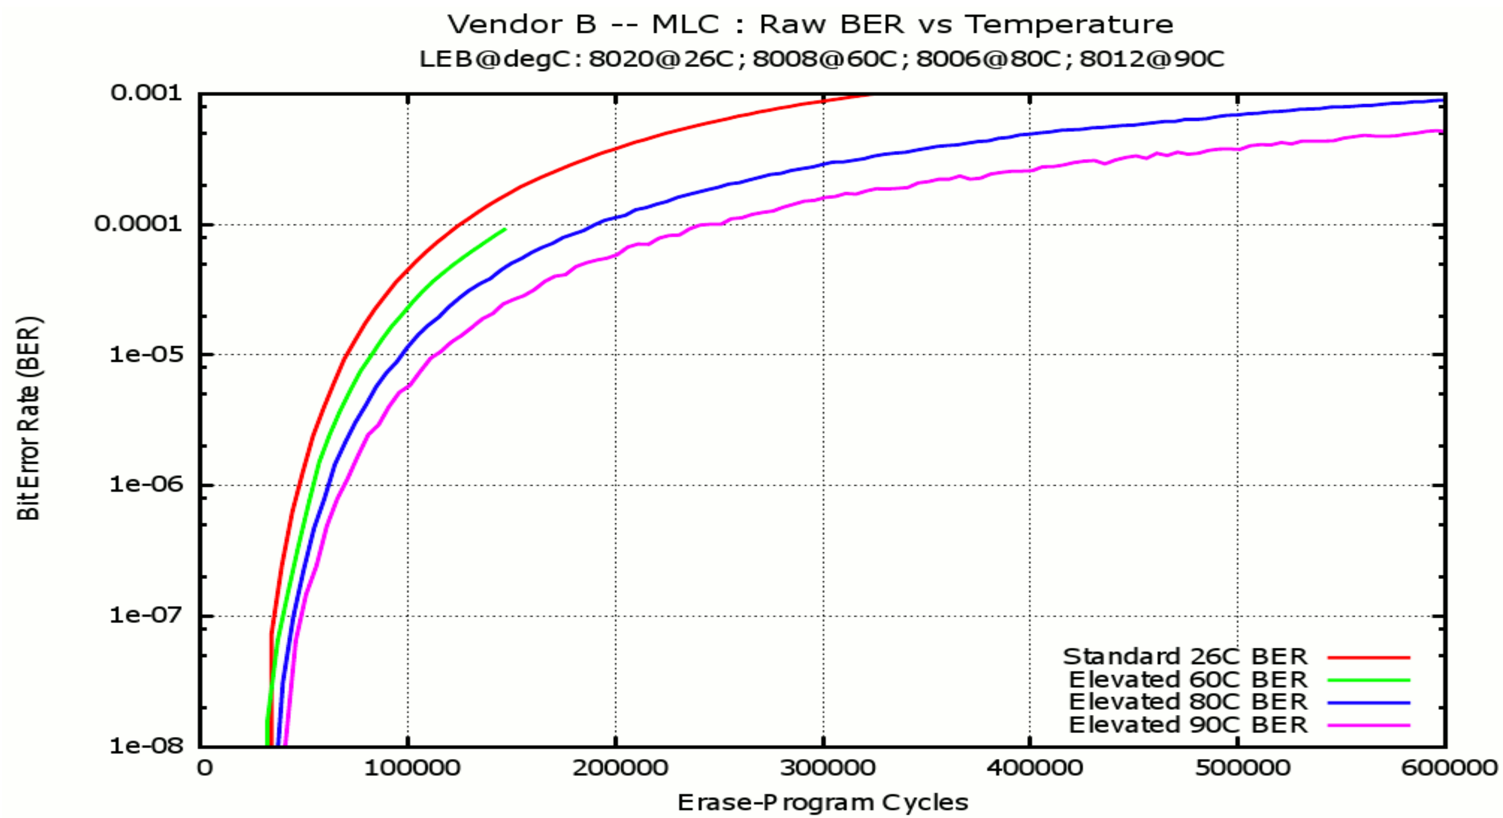

Error Rate Decreases with Increase in Temperature
But,
Data Retention Decreases with Temperature

Erase Cycle Effect on Read Disturbs

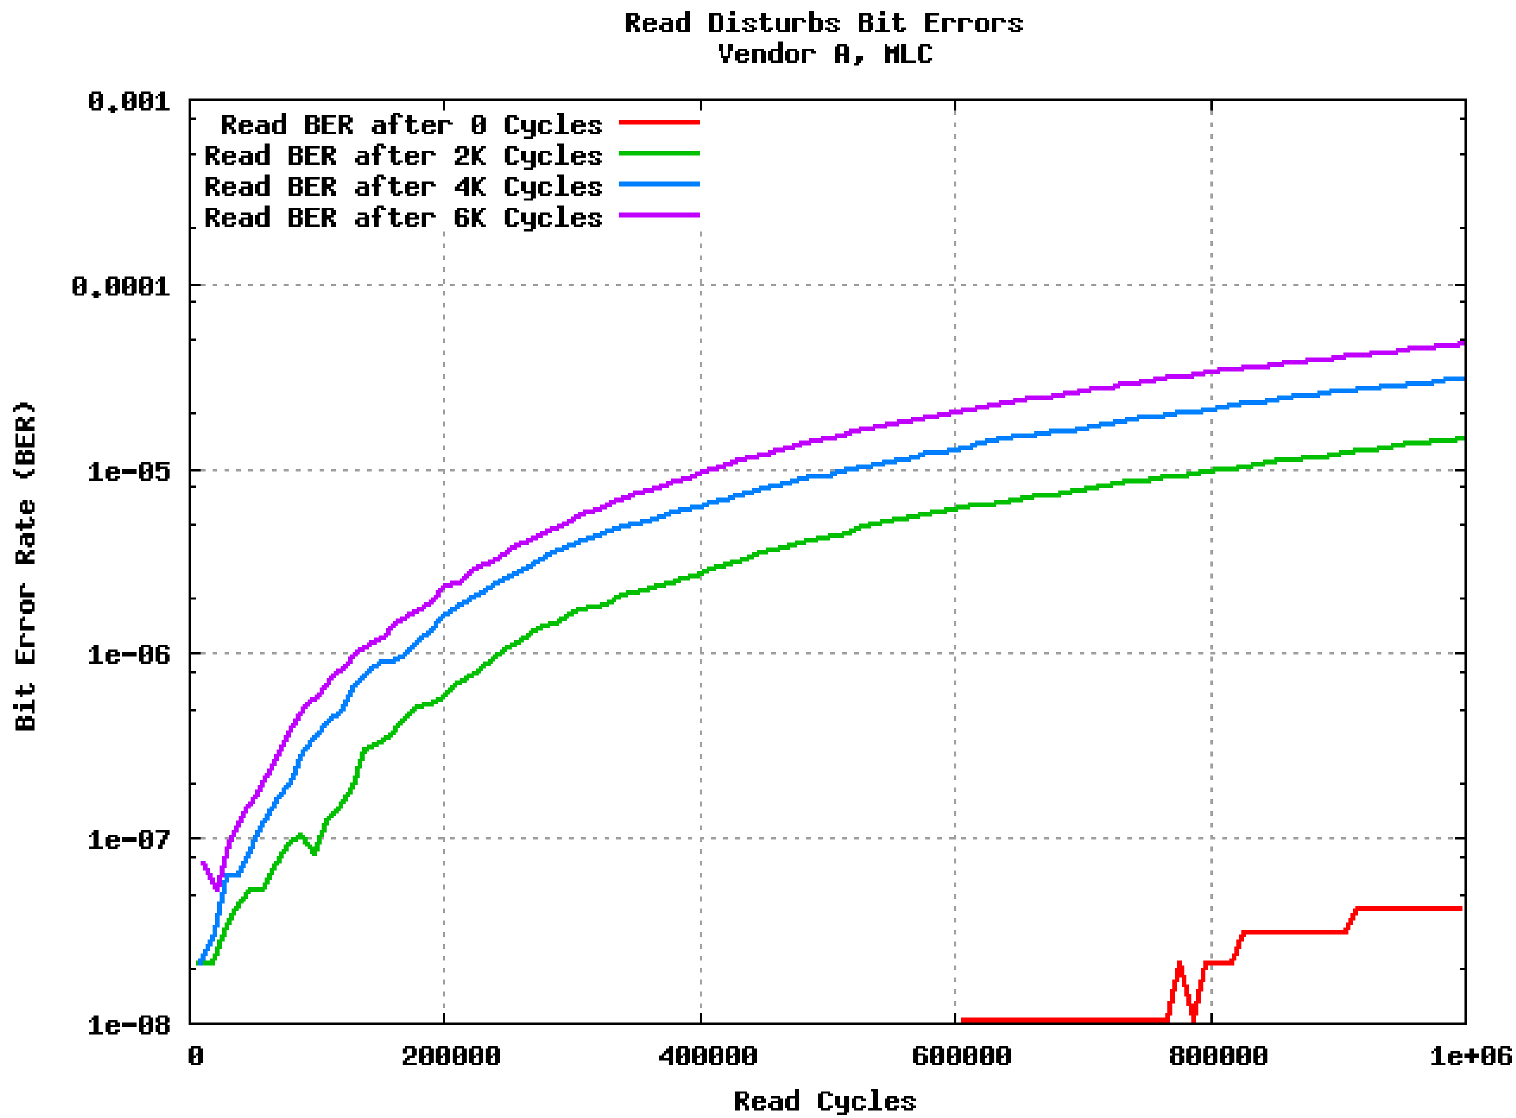

The Need For Robust ECC

The reason we have SSS is:

Thumb drives

Digital cameras

iPODs

Etc.

Drive the price down and capacity up.

These applications do not need:

Robust data integrity

Large number of write cycles

ECC Extends the PEB Life (SLC)

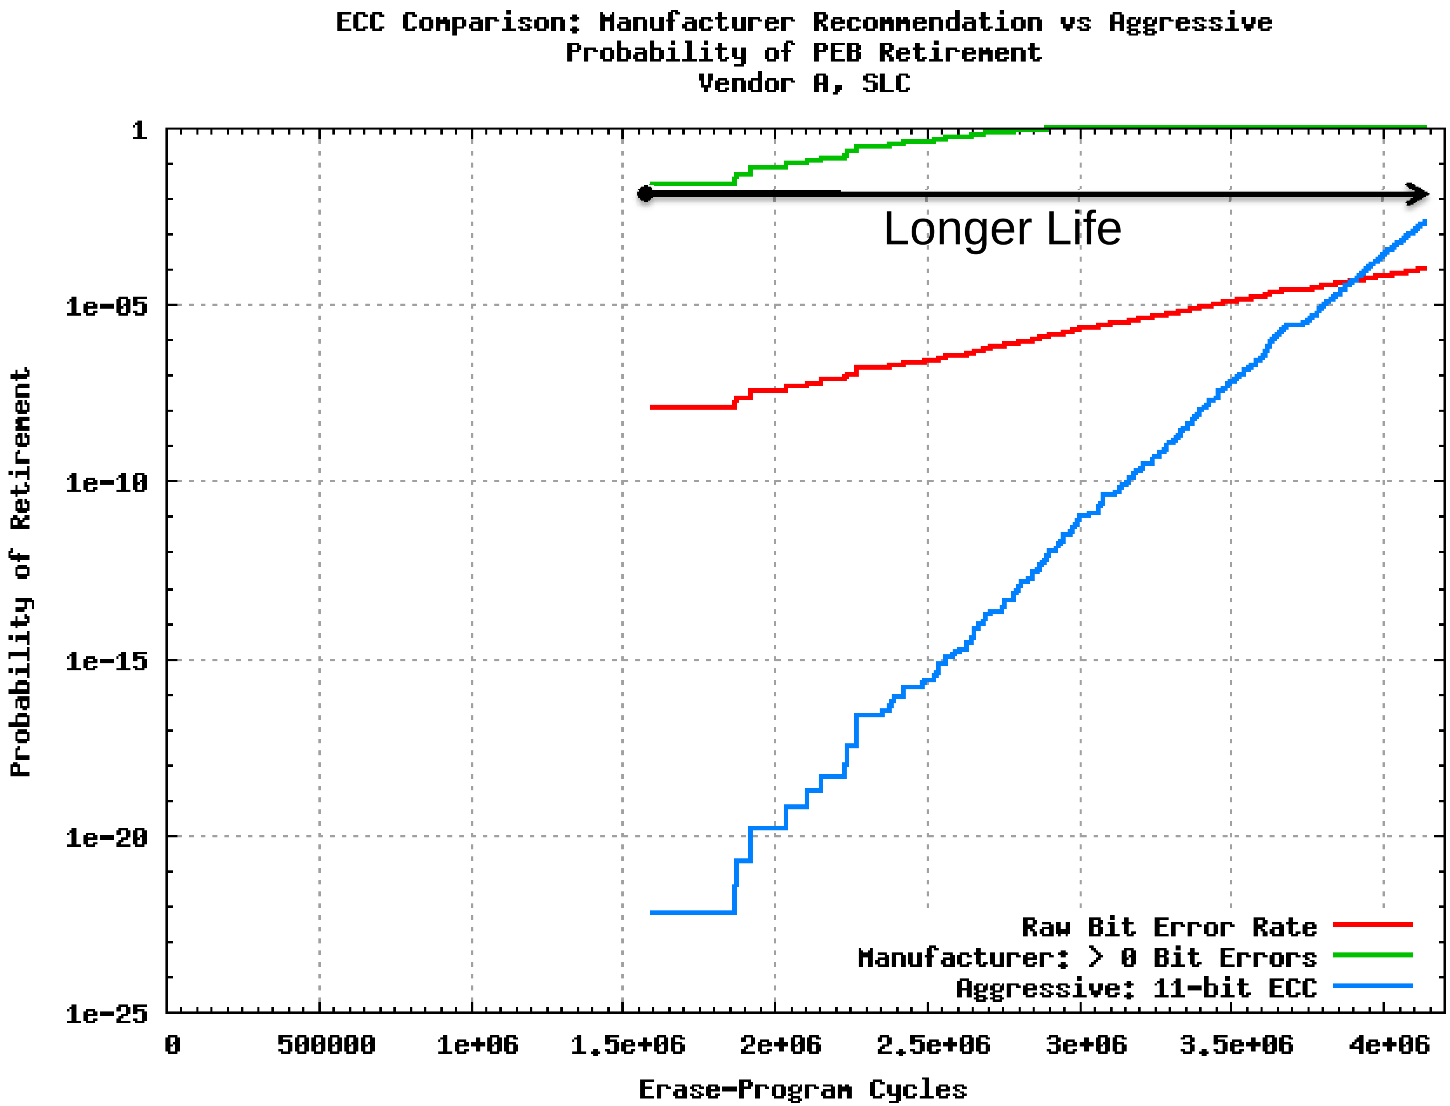

ECC Extends the PEB Life (MLC)

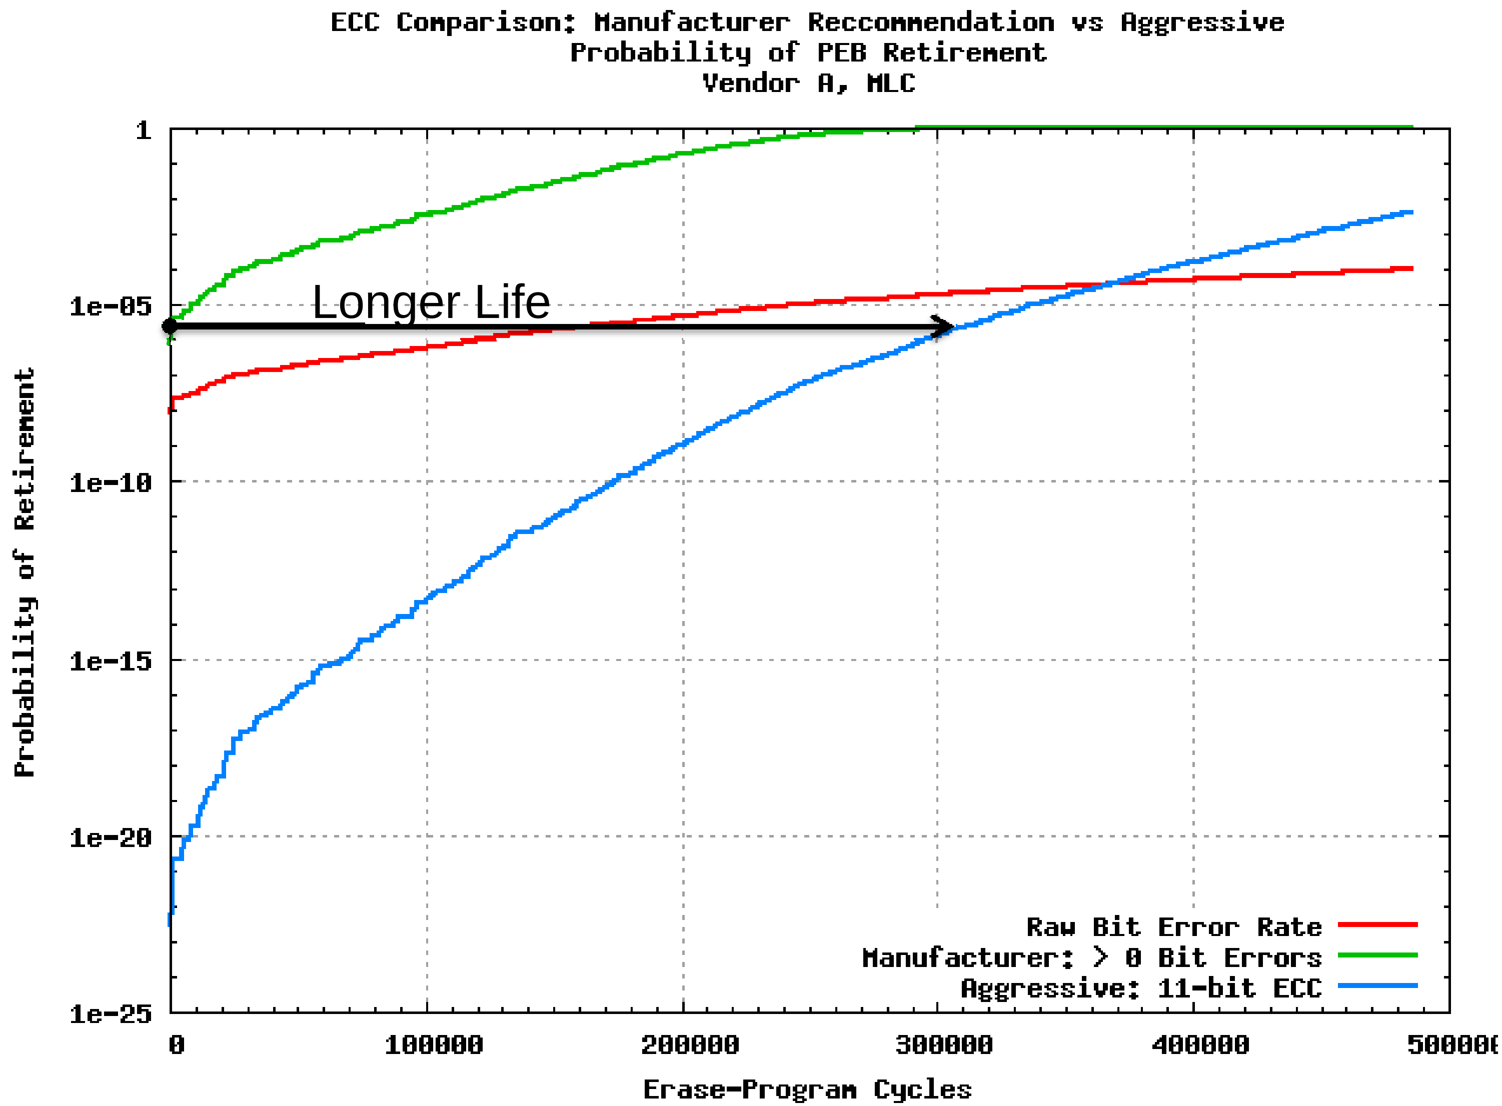

PEB Retirement v Temp. (MLC)

Education SNIA

MFG Recommendation | **Robust ECC / Aggressive Retirement**

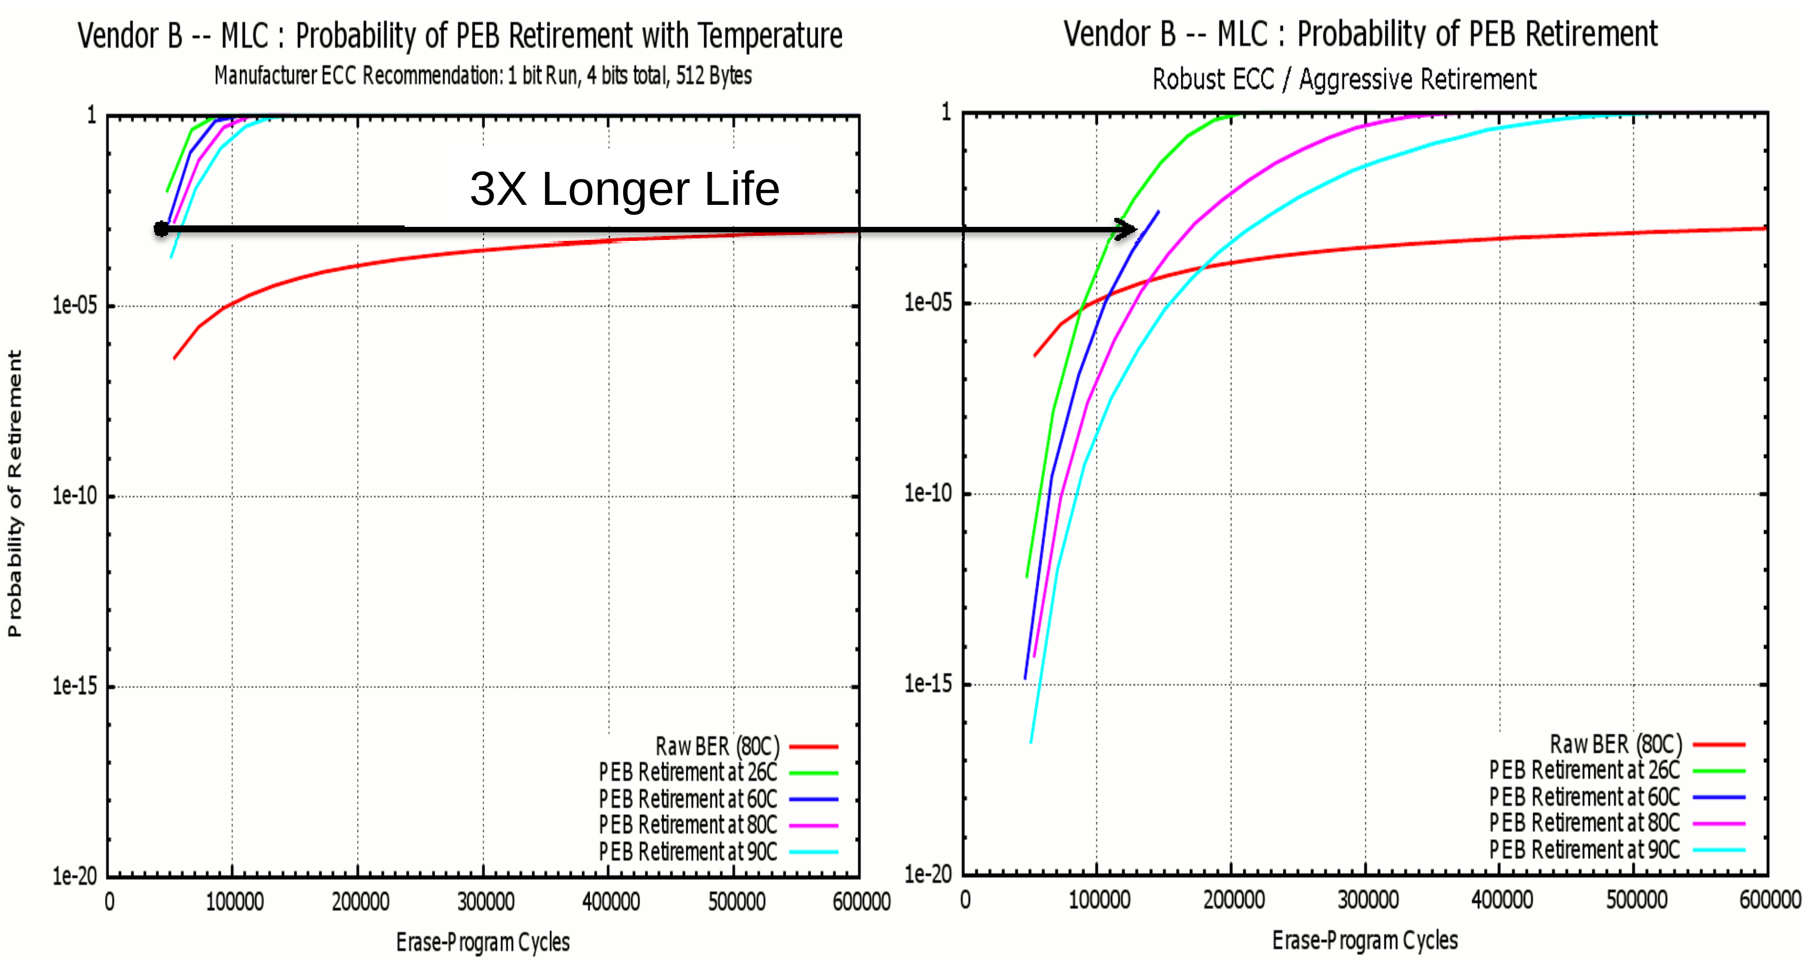

Probability of Exceeding ECC (MLC)

Education SNIA

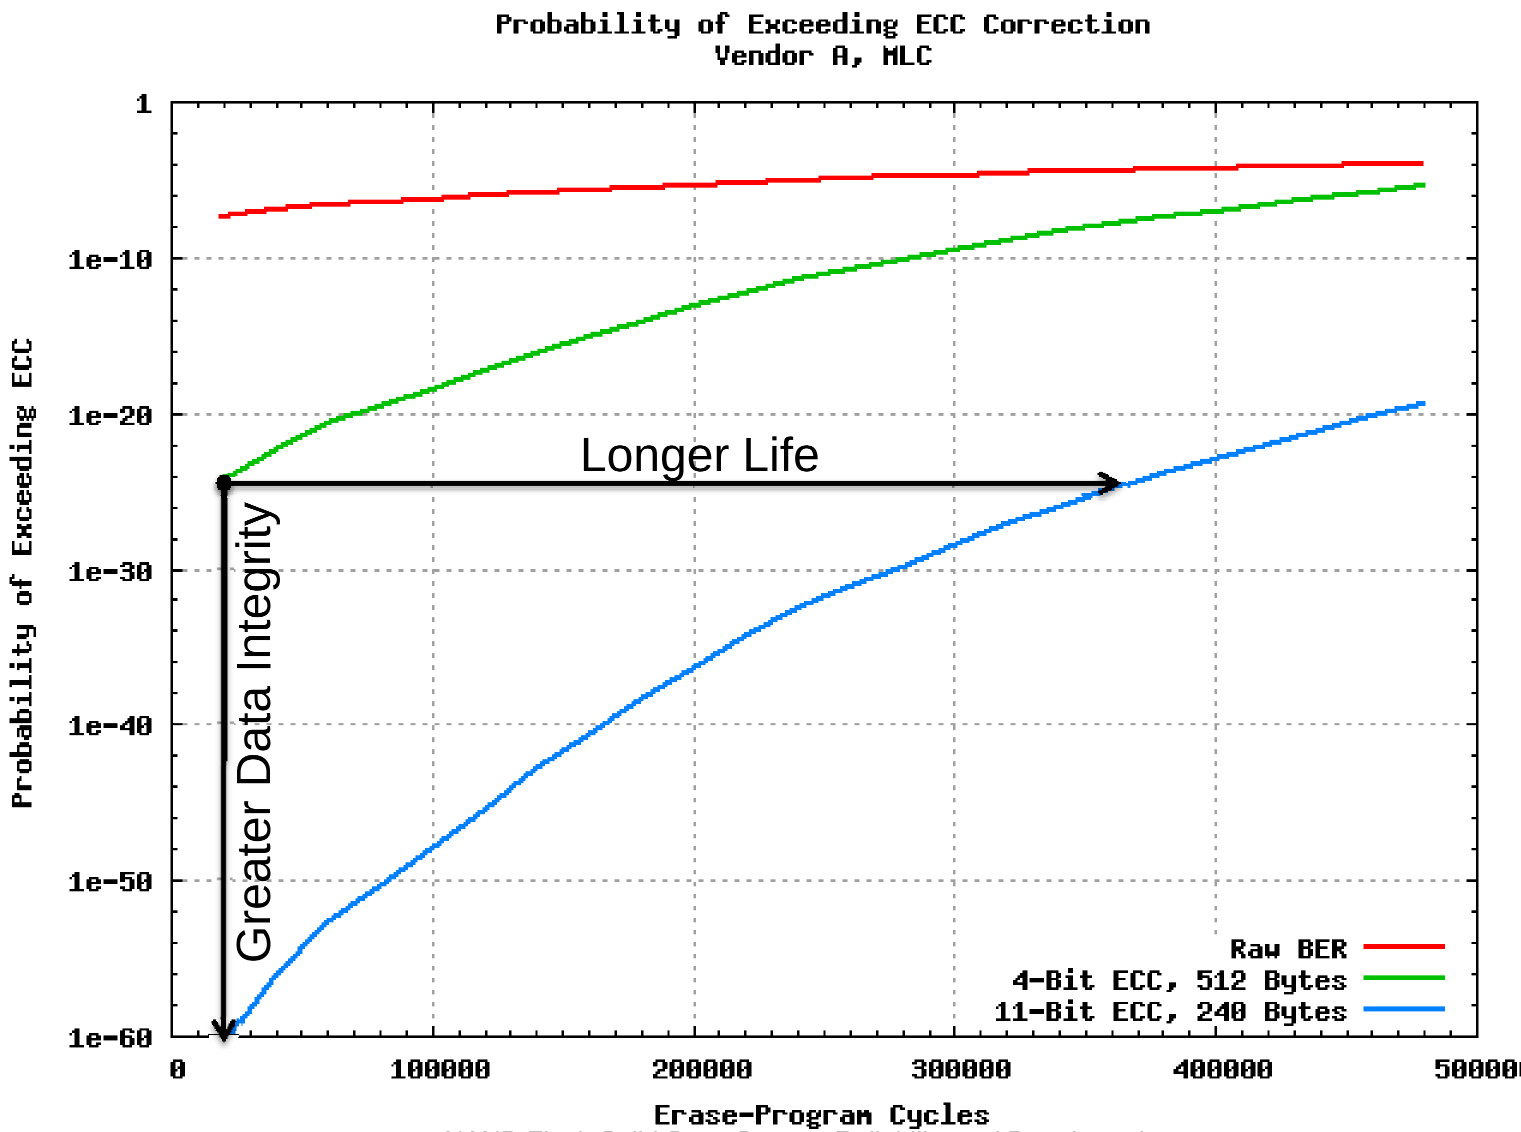

Data Integrity .V. Extended Life

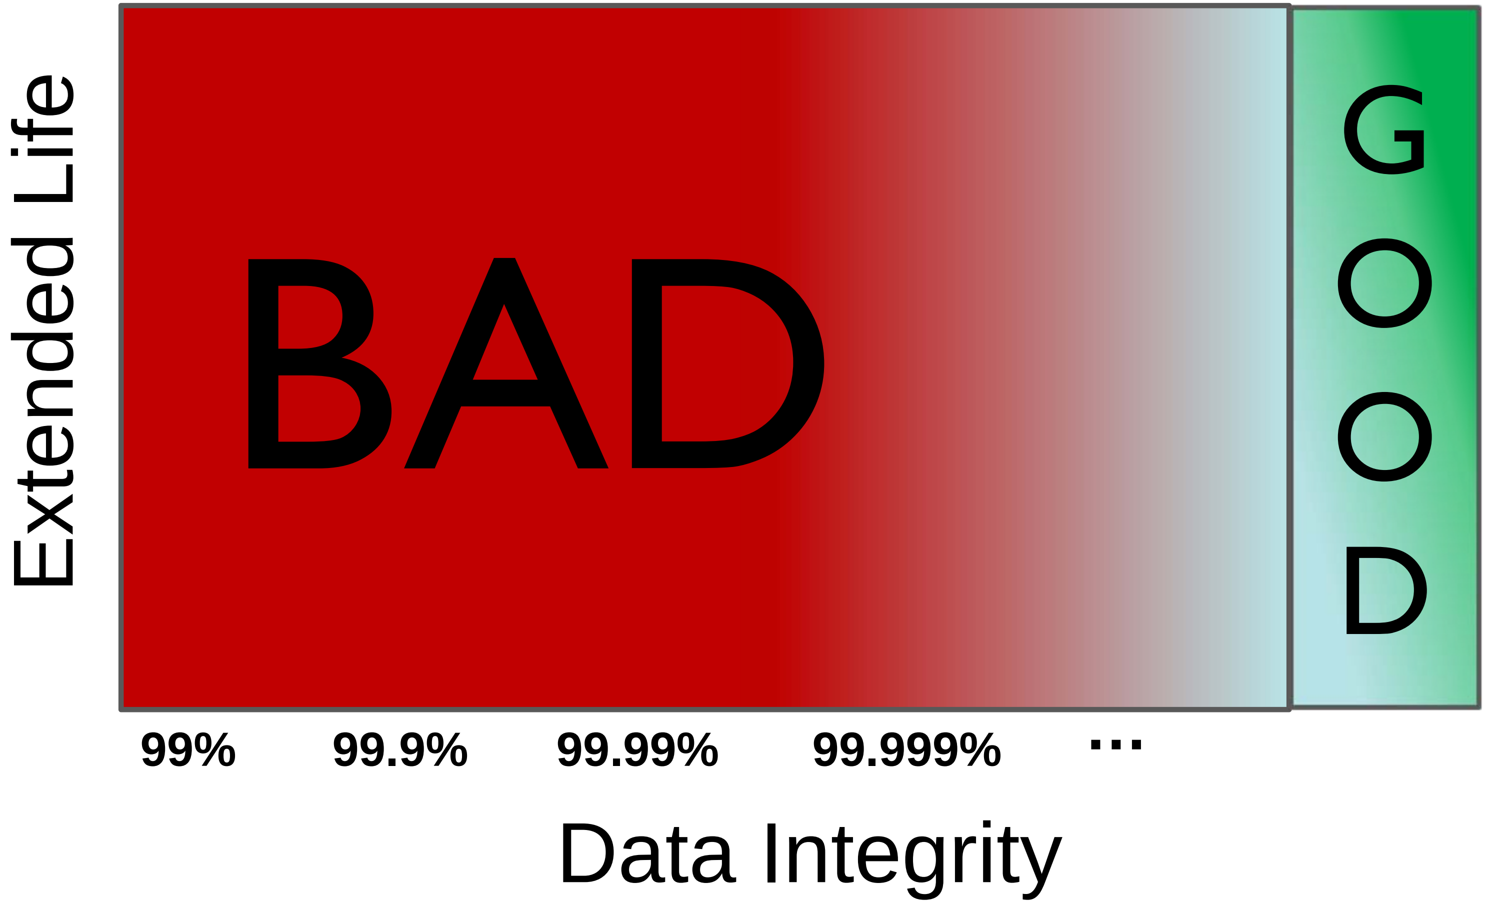

Life ("TBW : Terabytes Written")

	Life Expectancy	Average Rate	Used per Day	Consumed per Year	Life Expectancy
Consumer Vehicle	200K miles	45 MPH	1.5 Hrs	25K Miles	8.1 years
Commercial Vehicle	5,000K miles	55 MPH	11 Hrs	221K Miles	22.6 years
Consumer (MLC)	32,000 TBW	100 MB/s	7 Hrs	920 TB	34.8 years
Enterprise (SLC)	160,000 TBW	250 MB/s	24 Hrs	7,884 TB	20.3 years

Your Mileage May Vary

Life Calculation

"Enterprise"		"Consumer"
1,000,000	Program-Erase Cycles	100,000
160,000,000,000	Capacity (B)	320,000,000,000
160,000,000,000,000,000	Total Bytes Written	**32,000,000,000,000,000**
24	Hours used per day	7
86,400	Seconds used per day	25,200
250,000,000	Write Rate (B/s)	100,000,000
21,600,000,000,000	Bytes Written per day	2,520,000,000,000
7,884,000,000,000,000	Bytes Written per yr (B/yr)	**919,800,000,000,000**
20.3	Life in years	**34.8**

Access Process (Physics Ignored)

Read Access

- Address Chip / EB / Page
- Load Page into Register
- Transfer Data From Register 1-byte per cycle

Typical NAND Flash Die:
- 2000 Erase Blocks (EB)
- 64 Pages per EB
- 4000 Bytes per Page
- 500 MByte Total Capacity

Write Access

- Address Chip / EB
- Erase EB

…some time later…

- Address Chip / EB / Page
- Transfer Data To Register 1-byte per cycle
- Program Register to Page

Example 1: Read/Erase/Modify/Write

Education SNIA

Time = t1

Starting State

Page	Erase Block 1			
0	b	c	--	--
1	j	--	k	l
2	m	--	--	--
3	--	--	q	r

Time = t2

Write Buffer & W,X,Y

Page	Erase Block 1			
0	b	c	W	X
1	j	Y	k	l
2	m			
3			q	r

Time = t3

Write Buffer & Z,A,B',C',R'

Page	Erase Block 1			
0	B'	C'	w	x
1	j	y	k	l
2	m	Z	A	
3			q	R'

Buffer holds data while EB-1 Erased

Page	Erase Block 1			
0				
1				
2				
3				

Buffer holds data while EB-1 Erased

Page	Erase Block 1			
0				
1				
2				
3				

Explanation of Previous Slide

- Assumptions
  - Simplified to show erase blocks with 4 pages, each page having 4 data blocks
  - Invalid (erased or replaced) data is indicated by "—"
  - Old data is indicated by lower case letters
  - New data is indicated by CAPs; Replacement data is indicated by "prime" (e.g. c → C')
- Detail T = t1 to T = t2 transition
  - Data is read from EB-1
  - EB-1 is erased
  - New data {W,X,Y} modifies previous invalid data
  - Data is written back to EB-1
- Detail T = t2 to T = t3 transition
  - Data is read from EB-1 into data buffer
  - EB-1 is erased
  - New data {B', C', Z, A, R'} modifies previous data in data buffer
  - Data is written back to EB-1

Note: backup material for those reviewing or looking at presentation without audio/video

Example 2: Read/Modify/Write

Education SNIA

Time = t1
Starting State

Page	Erase Block 1			
0	b	c	--	--
1	j	--	k	l
2	m	--	--	--
3	--	--	q	r

Time = t2
Data to Buffer (not shown)
Erase EB-1 (not shown)
Write Buffer & W,X,Y to EB-1

Page	Erase Block 2			
0	b	c	W	X
1	j	Y	k	l
2	m			
3			q	r

Time = t3
Data to Buffer (not shown)
Erase EB-1 (not shown)
Write Z,A & Replace b,c,r with B',C',R' & Write EB-1

Page	Erase Block 3			
0	B'	C'	w	x
1	j	y	k	l
2	m	Z	A	
3			q	R'

Implicit wear leveling; EB-1 → EB-2 → EB-3
Presumes that destination EB-2 & EB-3 erased prior to transfer of data → higher performance (than previous "Read/Erase/Modify/Write" example)

"Write Amplification Impact"

- In this example,
  - Data written t1 to t2: 16 blocks
    - NEW DATA {W, X, Y} 3 blocks; Copied Data {b, c, j, k, l, m, q, r} 8 blocks
    - Null Data: 5 blocks
  - Data written t2 to t3: 16 blocks
    - NEW DATA {B', C', Z, A, R'} 5 blocks; Copied Data {w, x, j, y, k, l, m, q} 8 blocks
    - Null Data: 3 blocks
  - (2) EB erasures
  - 25% (8 of 32) writes are user initiated
  - 75% (24 of 32) writes are internal data movement (overhead)
- Important:
  - Amount of valid or invalid data in EB-1 is irrelevant to performance impact
  - "Write Amplification" is workload (access pattern) dependent (e.g., what if the write of R' above was not coincident with B' & C')

SATA-B: IOPS vs Transfer Size

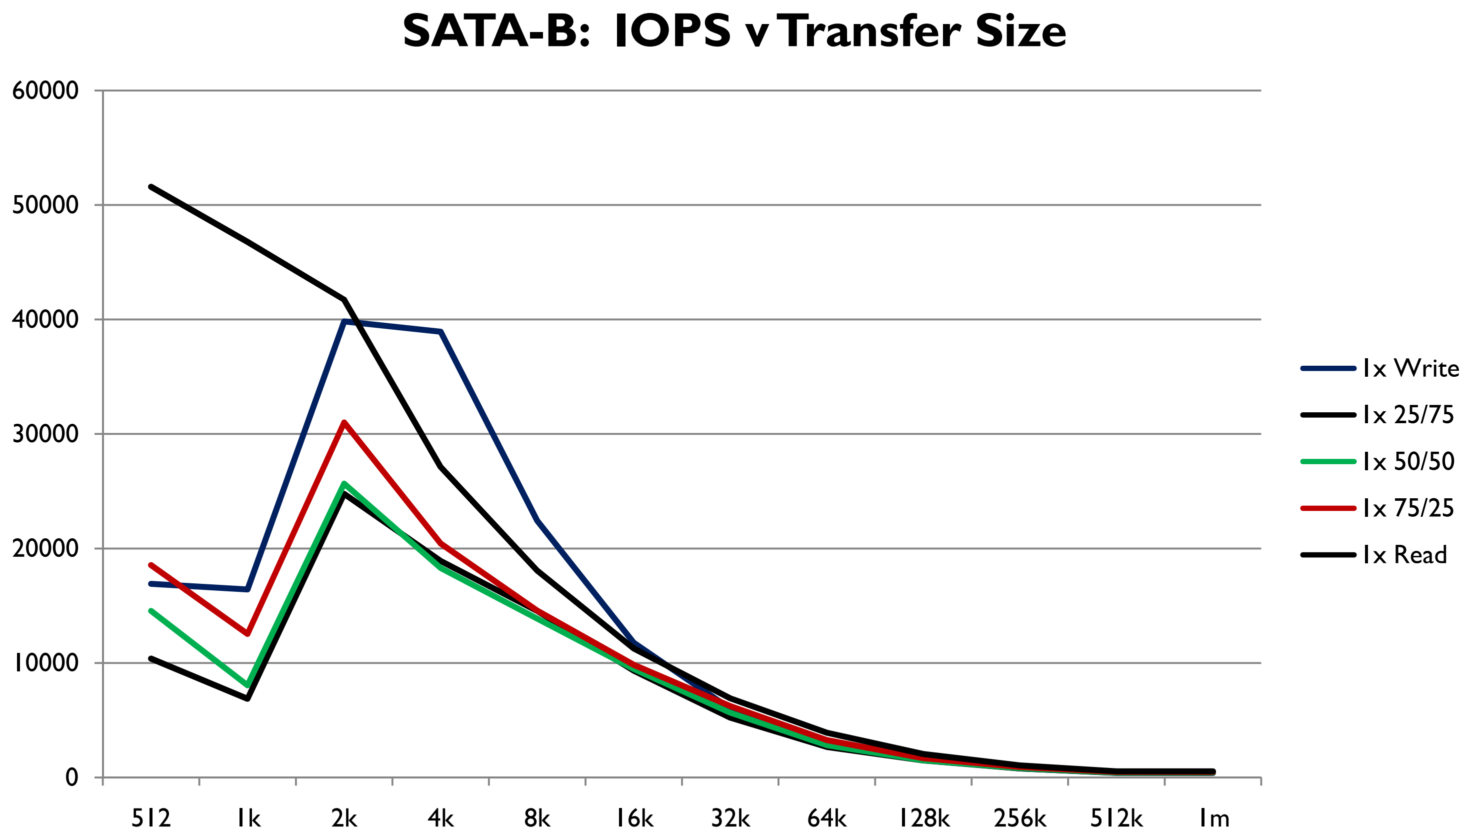

Possible evidence of "write amplification"

Example 3: Garbage Collection

Education SNIA

Time = t1

Start Garbage Collect EB-1

Page	Erase Block 1			
0	b	c	--	--
1	j	--	k	l
2	m	--	--	--
3	--	--	q	r

Page	Erase Block 2			
0				
1				
2				
3				

Time = t2

EB-1 GC'd to EB-2

W,X,Y added

Page	Erase Block 1			
0	b	c	--	--
1	j	--	k	l
2	m	--	--	--
3	--	--	q	r

Page	Erase Block 2			
0	W	b	c	X
1	Y	j	k	l
2	m	q	r	
3				

Time = t3

EB-1 erase

b,c,r replaced by B',C',R'

Page	Erase Block 1			
0				
1				
2				
3				

Page	Erase Block 2			
0	w	--	--	x
1	y	j	k	l
2	m	q	--	B'
3	C'	Z	A	R'

Explanation of Previous Slide

- Assumptions
  - New data blocks and data blocks being garbage collected are interleaved
- Details
  - At Time = t0, erase block 1 (EB-1) is identified for Garbage Collection (GC)
  - At Time = t1, good data is moved from EB-1 to EB-2 (it is implicit that an index is updated accordingly); New data W, X, and Y are added while the GC is taking place. EB-1 is then ready to be erased
  - At Time = t2, EB-2 is erased; Data for b, c, & r have been updated with B', C' & R'; b, c & r are indicated as "Invalid."

Note: backup material for those reviewing or looking at presentation without audio/video

GC Performance Impact

Education SNIA

- In this example,
  - COPIED DATA: {b, c, j, k, l, m, q, r} 8 blocks
  - NEW DATA {W, X, Y, B', C', Z, A, R'} 8 blocks
  - 50% (8 of 16) writes are user initiated
  - 50% (8 of 16) writes are internal movement (overhead)
- Important:
  - 50% of EB-1 was "invalid data"
  - What if only 10% had been "invalid data?"
  - GC efficiency is dependent upon % of reserve capacity

GC: Pathological Write Conditions

- IF high percentage of total storage capacity utilized

AND

- High percentage of data has no correlation-in-time

AND

- Continuous writing (no recovery time for GC)

THEN

Efficiency of GC greatly diminished

Pathological Write Condition

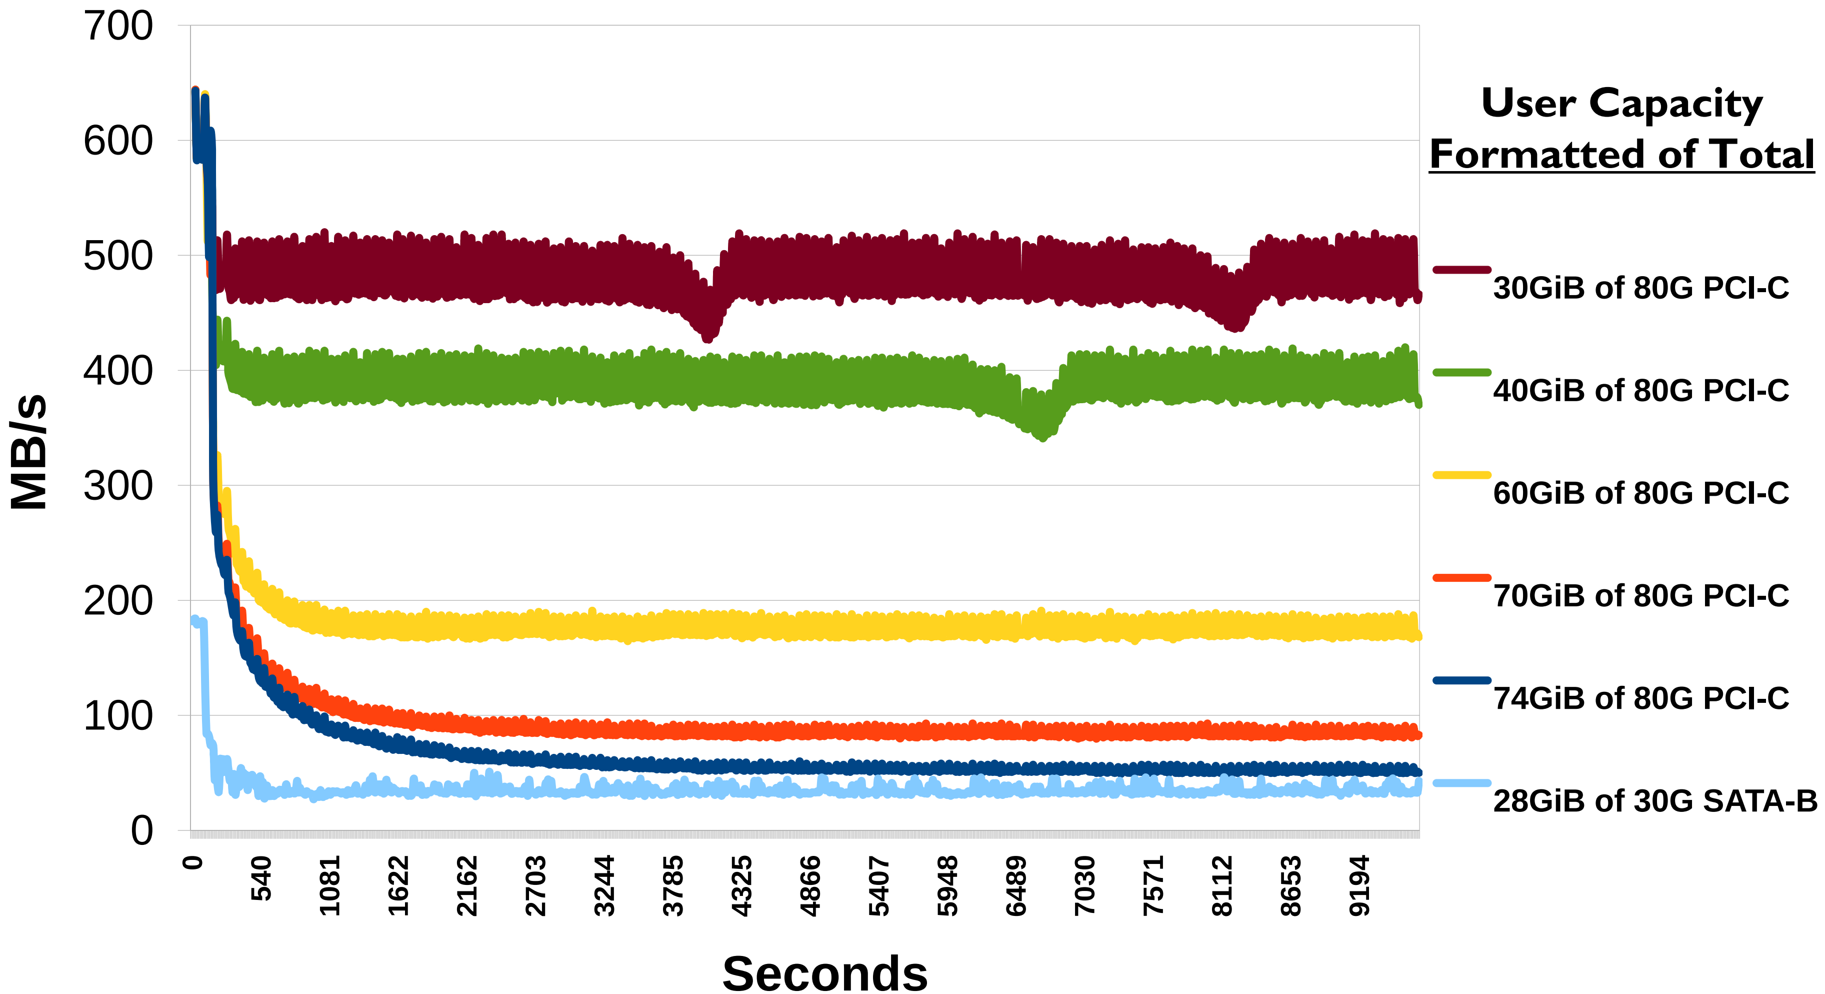

Your Mileage Will Vary

Modern Enterprise SSS Controllers

- SSS Solutions have these features now

MANY DO NOT

- Do your Homework
- Choose Wisely
- Caveat Emptor

SSS Data Protection Mechanisms

Education SNIA

- Robust ECC
  - Increased need for MLC
  - Increased need as NAND Flash Capacity Increases
- Index Protection (ECC, DIF, Other)
  - Eliminates Index Induced Silent Data Corruption
- Internal RAID
  - Overcomes High Infant Mortality
  - Continued Protection Over Life

SSS Data Protection Mechanisms

- Power Loss
  - Storage of metadata / indexes
  - Completion of acknowledged writes
- Data Scrubbing
- Read / Program Disturb Management
- Data Retention Management
- Thermal Management
- Partial Page Programming Avoidance

SSS Life Extension Mechanisms

- Robust ECC
- RAID (Internal)
- Wear Leveling – Reduce Hot Spots
- Wear-out Prediction – Planned Maintenance
- Write Amplification Avoidance
- Garbage Collector Efficiency

Reliability – External Conditions

- System Configuration
  - External RAID
  - Write Caching
  - DIF Support
- Application Characteristics
  - Read/Write Ratios
  - Temporal Randomness of Access
  - Steady State v Bursting
- Configuration / Tuning
  - Reserve Capacity Setting (% of used capacity)
- Device Management
  - Device Retirement (device wear affects error probabilities)

Contra-indicated?

- Bandwidth Throttling
  - May improve device life
  - May not translate to the system / data center or more global reliability gains.
- Partial Page Programming
  - May decrease write amplification
  - Severely reduces data integrity
- SSS tuned for performance at a block size you don’t use
  - Write amplification at block size you do use
  - Increased wear
  - Decreased data integrity
  - Reduced life

Global v Local Optimization

Education SNIA

- **Compare at data center / system level**
  - Performance
  - Cost (CapEx; OpEx)
  - Power
  - **Reliability, Availability**
- **Tune for Your Needs!**

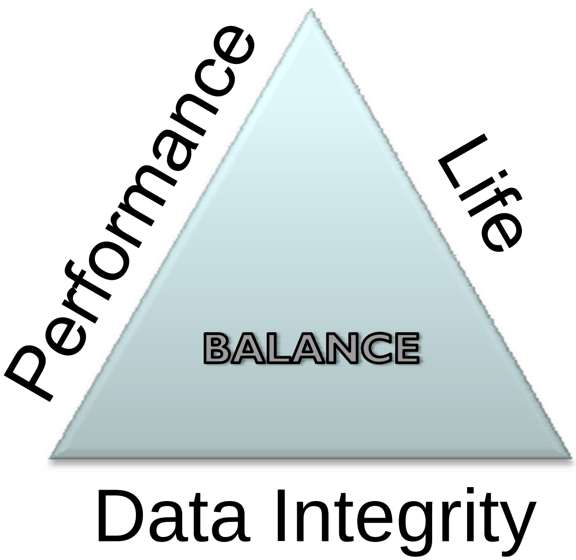

Q&A / Feedback

Education SNIA

- Please send any questions or comments on this presentation to SNIA: ***tracksolidstate@snia.org***

Many thanks to the following individuals for their contributions to this tutorial.

- SNIA Education Committee

Khaled Amer
Rob Peglar
Marius Tudor