



Enterprise MLC NAND Industry Comparison

Gary Tressler, Dustin Vanstee and Tom Griffin
IBM Corporation



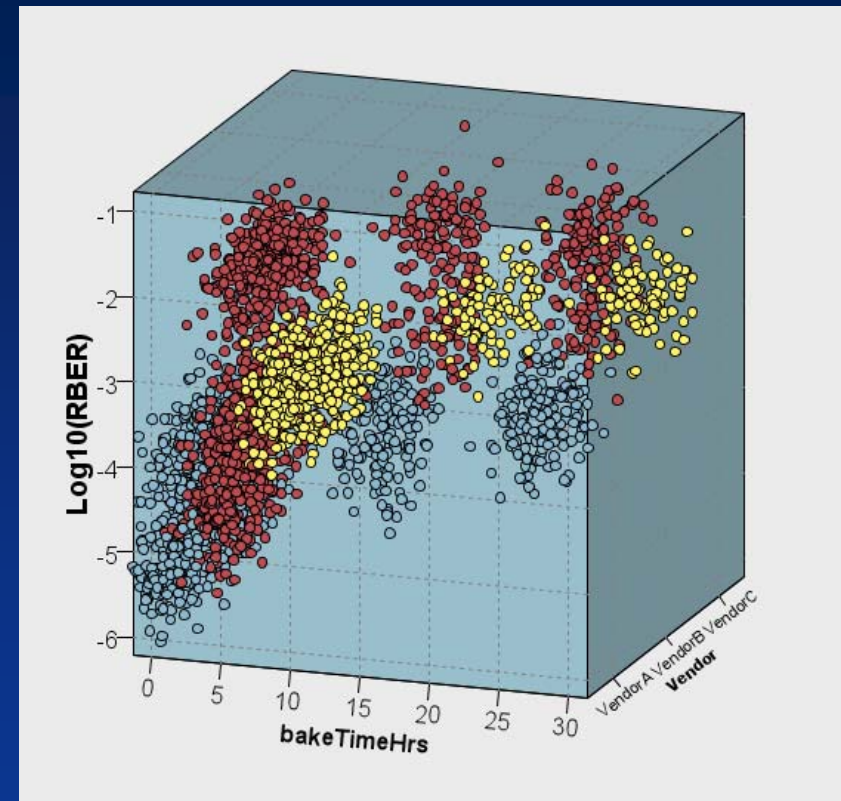
Agenda

- Introduction
- Goals
- Methodology
- Platform
- Program-Erase Cycling
- High Temperature Data Retention
- Summary

Flash Characterization Goals

Goals

- Compare industry Enterprise MLC NAND devices
 - Multiple suppliers
 - Endurance/retention BER envelope
 - Program-erase cycling temperature sensitivity
 - Dwell time sensitivity
 - UBER analysis
- Summarize relationship of characterization variables relative to Raw Bit Error Rate (RBER)
- Attempt to understand Flash impacts on SSD usable life
 - Use characterization data to interpolate for different usage scenarios
 - Use characterization data to extrapolate to usage scenarios that are time prohibitive to directly measure



Example of RBER variation across Flash blocks

Why Does Flash Bit Error Rate Matter for SSDs?

- In enterprise applications, Flash is generally written at a much higher rate than in the client space
- Flash cells degrade due to the large voltage required to program/erase the devices and due to the presence of defects
 - After many program/erase cycles it is not always possible to read back the data stored due to physical wearout of the device cell
- To overcome this effect, SSD controllers implement various ECC and recovery mechanisms to mitigate bit errors
 - ECC and recovery schemes can only protect to a limit
- From an SSD perspective, validating this limit is a time consuming process – will likely take years to understand without acceleration
- Flash characterization requires testing under different environmental and usage conditions, and gathering Raw Bit Error Rate (RBER) statistics to evaluate the Flash robustness



Flash Characterization Methodology

Enterprise MLC NAND
Industry Comparison

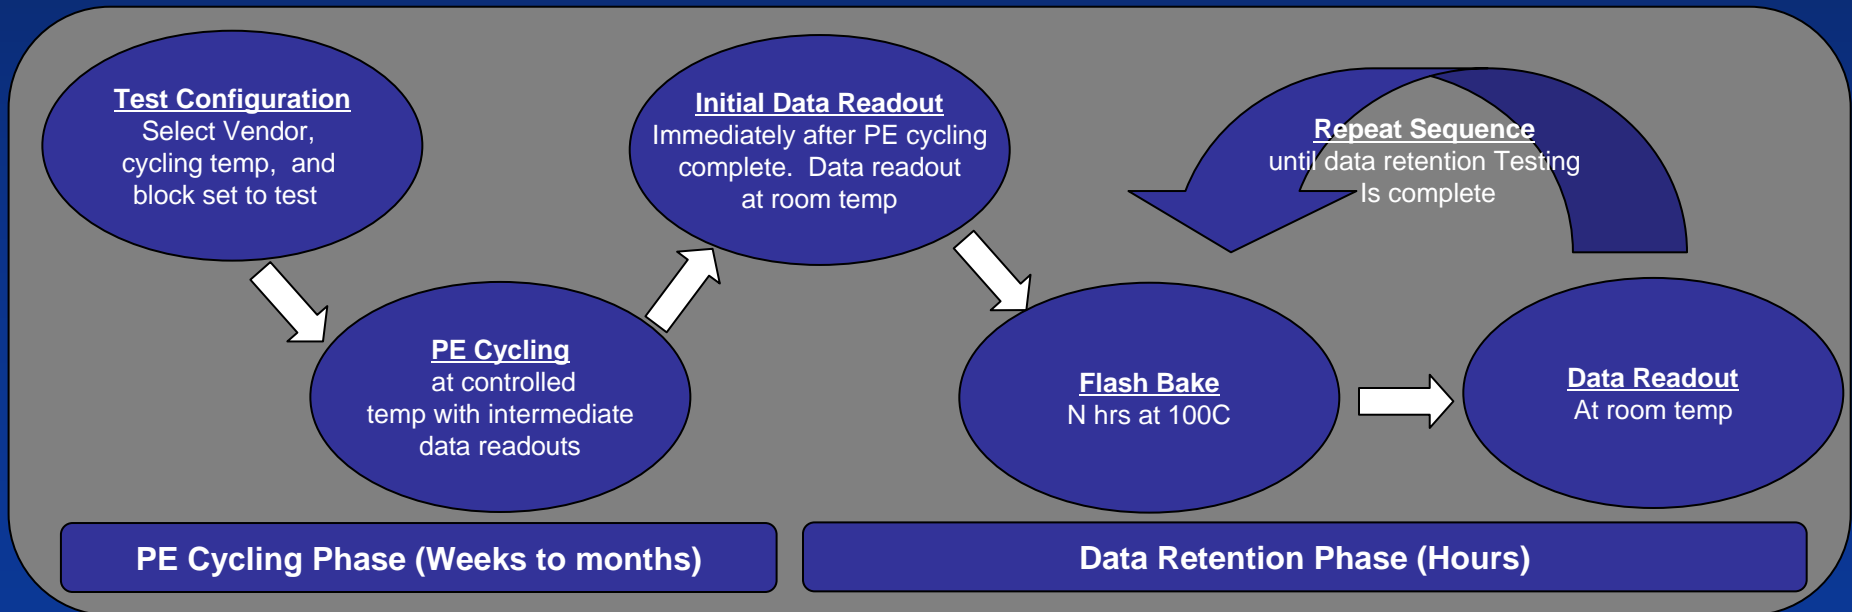
Approach

- Dwell time = 30, 60s, 120s, 240s, 480s
- Program-erase cycling = 15K, 30K
- Cycling temperature = 25C, 55C, 85C
- Bake temperature = 100C
- Bake hours = 0, 1, 2, 3, 4, 15, 26 hrs

PE Cycling Data Set Matrix			
	Vendor		
PeCyclingTemp	A	B	C
25	30,60,120,240,480	30,60,120,240,480	30,60,120
55	30,60,120,240,480	30,60,120,240,480	x
85	30,60,120,240,480	30,60,120,240,480	x

Data Retention Data Set Matrix			
	Vendor		
PeCyclingTemp	A	B	C
25	30,60,120	30,60,120	30,60,120
55	30,60,120	30,60,120	x
85	30,60,120	30,60,120	x

* Dwell times applied noted under each supplier



Flash Characterization Platform


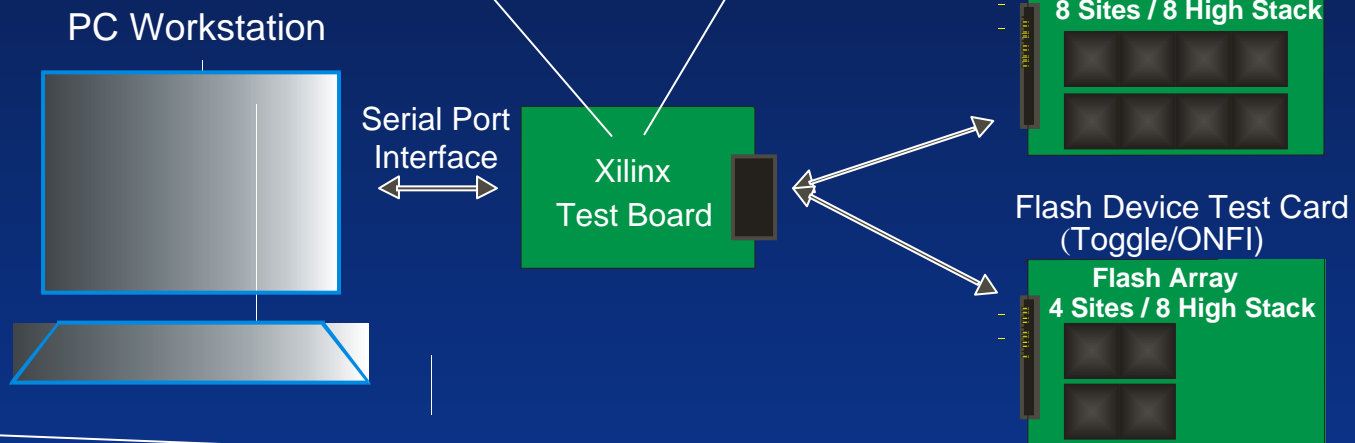
Enterprise MLC NAND
Industry Comparison

FPGA Evaluation Board


- Embedded processor
- Drives signaling and test data to Flash test cards
- Summarizes test responses and transmits via serial port to PC

Logic

- Cross-device sampling
 - Hides tPROG and tBERS and
 - Optimizes test card bandwidth
- Nested cycling permits multiple dwell times



25C, 55C, 85C PE Cycling

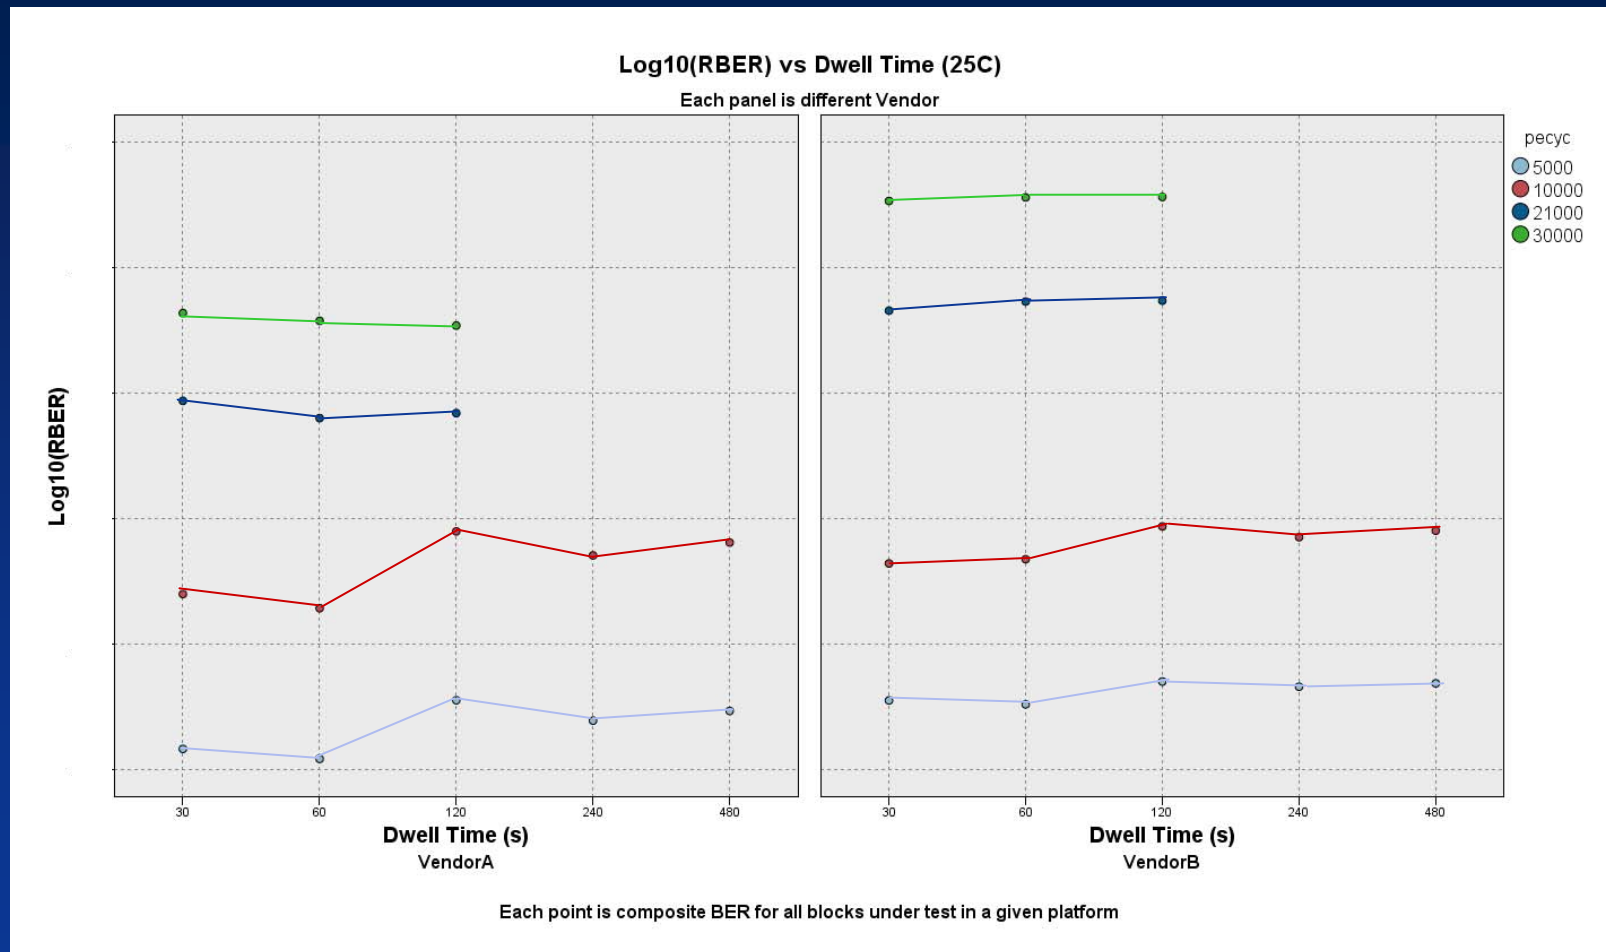




Raw Program-Erase Cycling @ 25C

No Data Retention, Dwell Time Variance

Enterprise MLC NAND
Industry Comparison

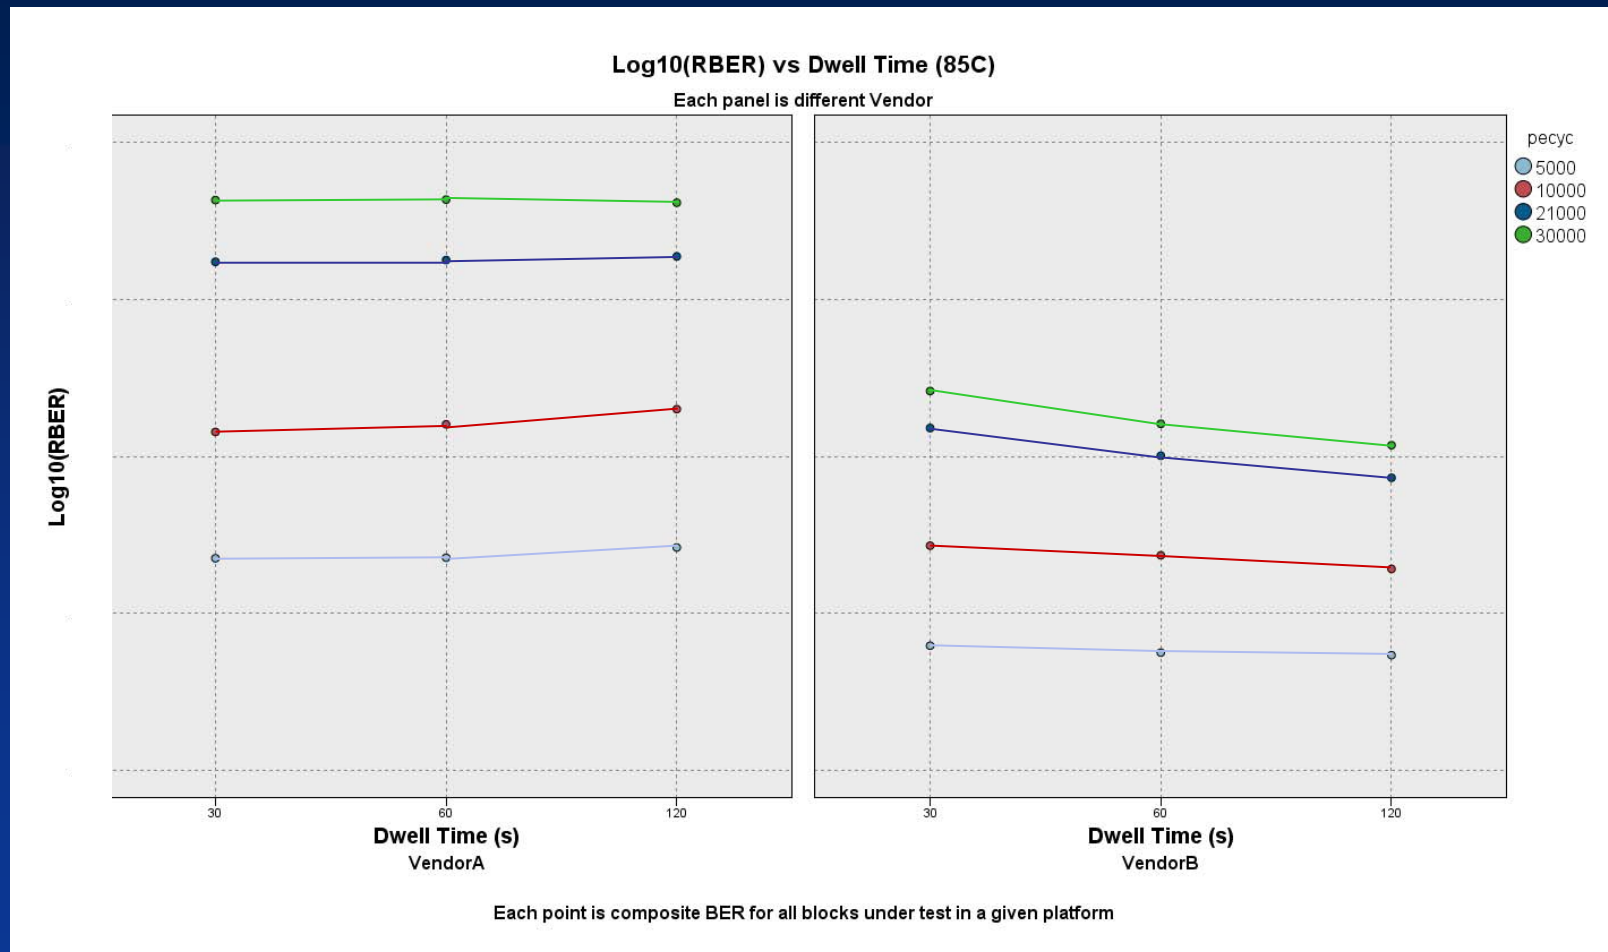


Raw program-erase cycling (no data retention evaluation) - RBER at 25C

- Small and inconsistent sensitivity to dwell time

Raw Program-Erase Cycling @ 85C

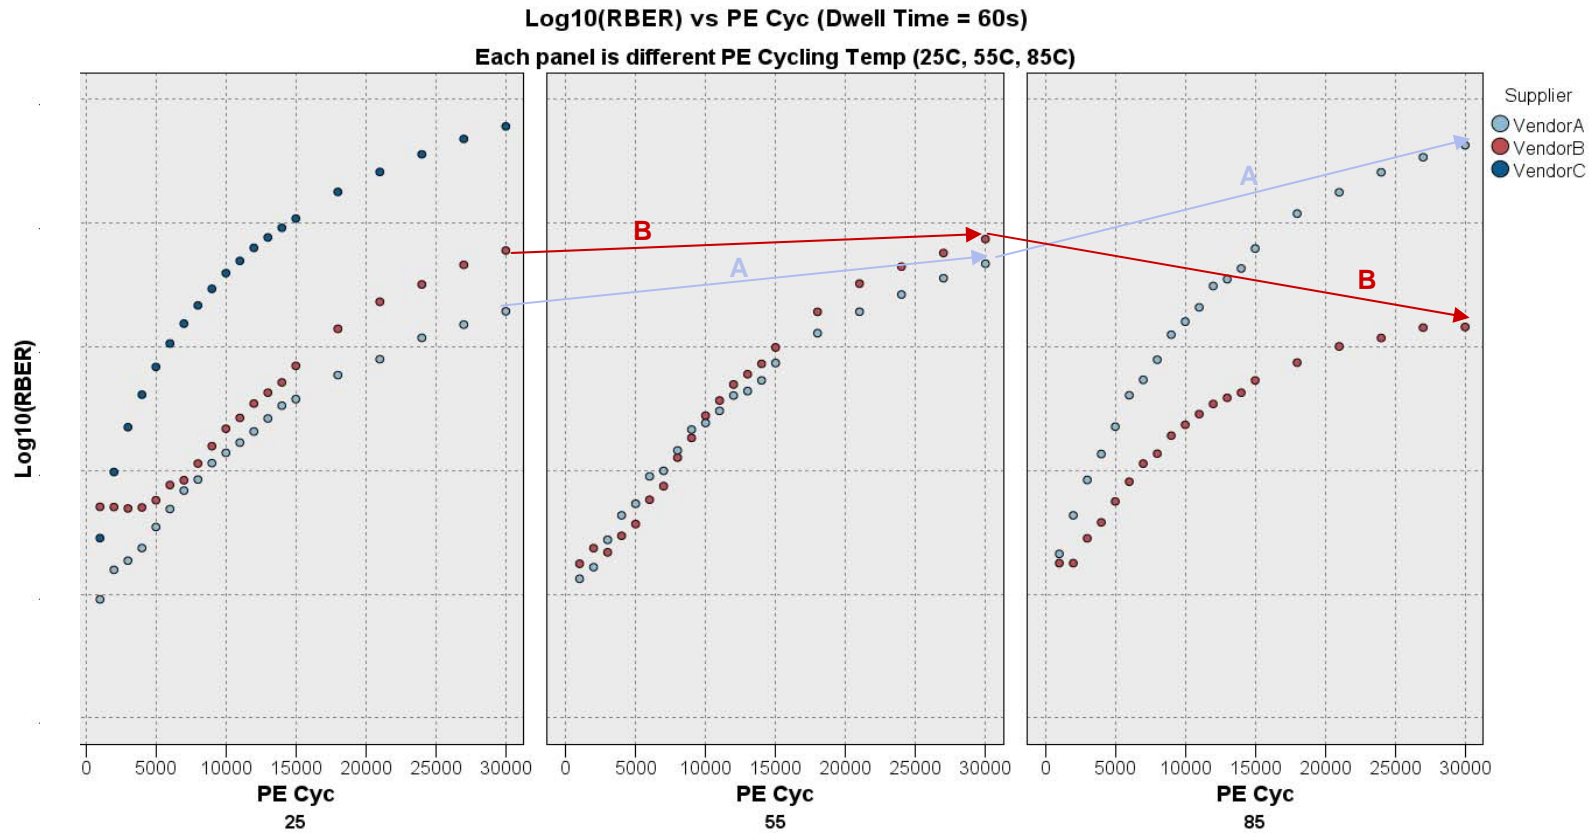
No Data Retention, Dwell Time Variance



Raw program-erase cycling (no data retention evaluation) - RBER at 85C

- Vendor A shows minimal sensitivity to dwell time
- Vendor B shows some evidence of longer dwell time resulting in lower RBER

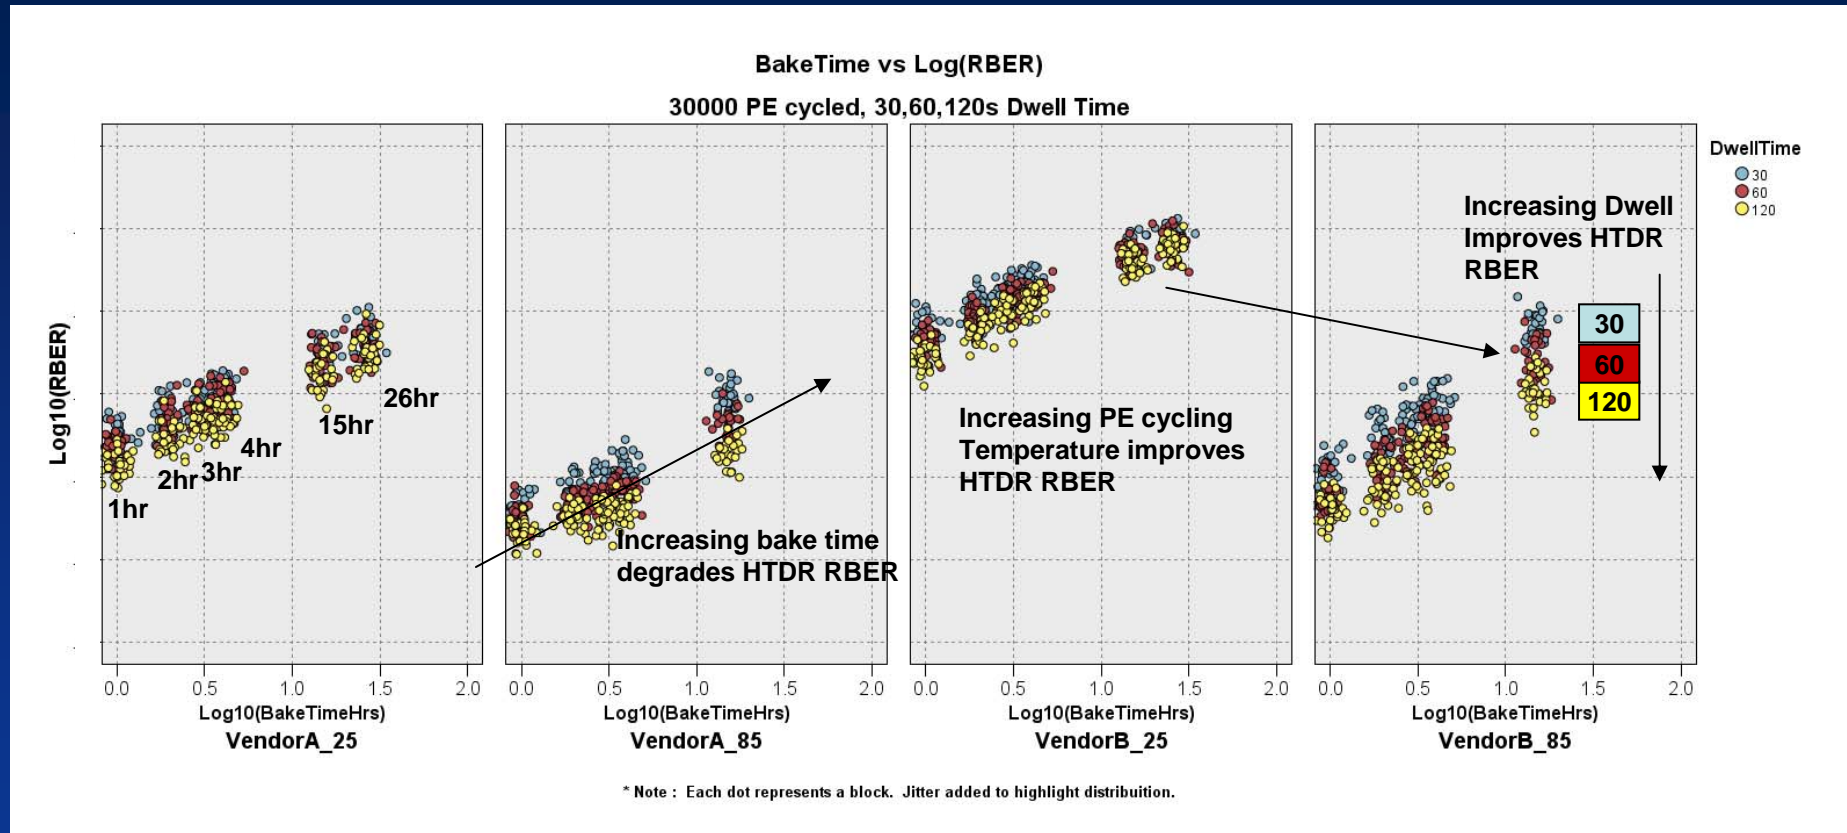
Raw Program-Erase Cycling No Data Retention, Temperature Variance



Each point is composite BER for all blocks under test in a given platform (vendor, temperature, dwell time)

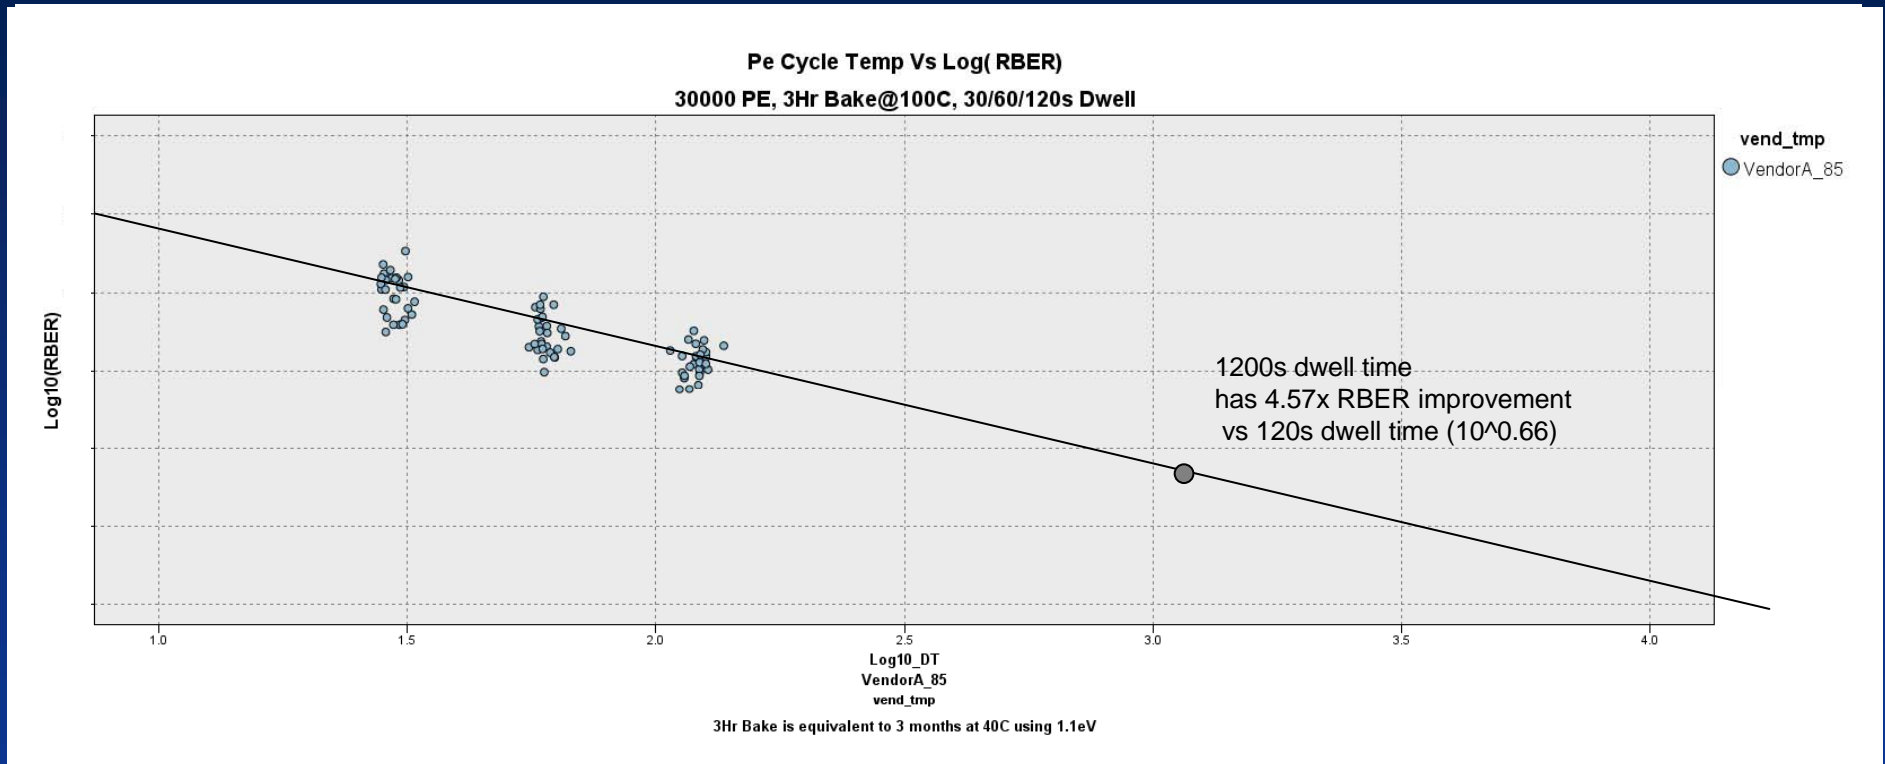
- Program-erase cycling temperature sensitivity inconclusive**
- Vendors A & B show opposite relationships relative to temperature

Post Program-Erase Cycling Data Retention Study



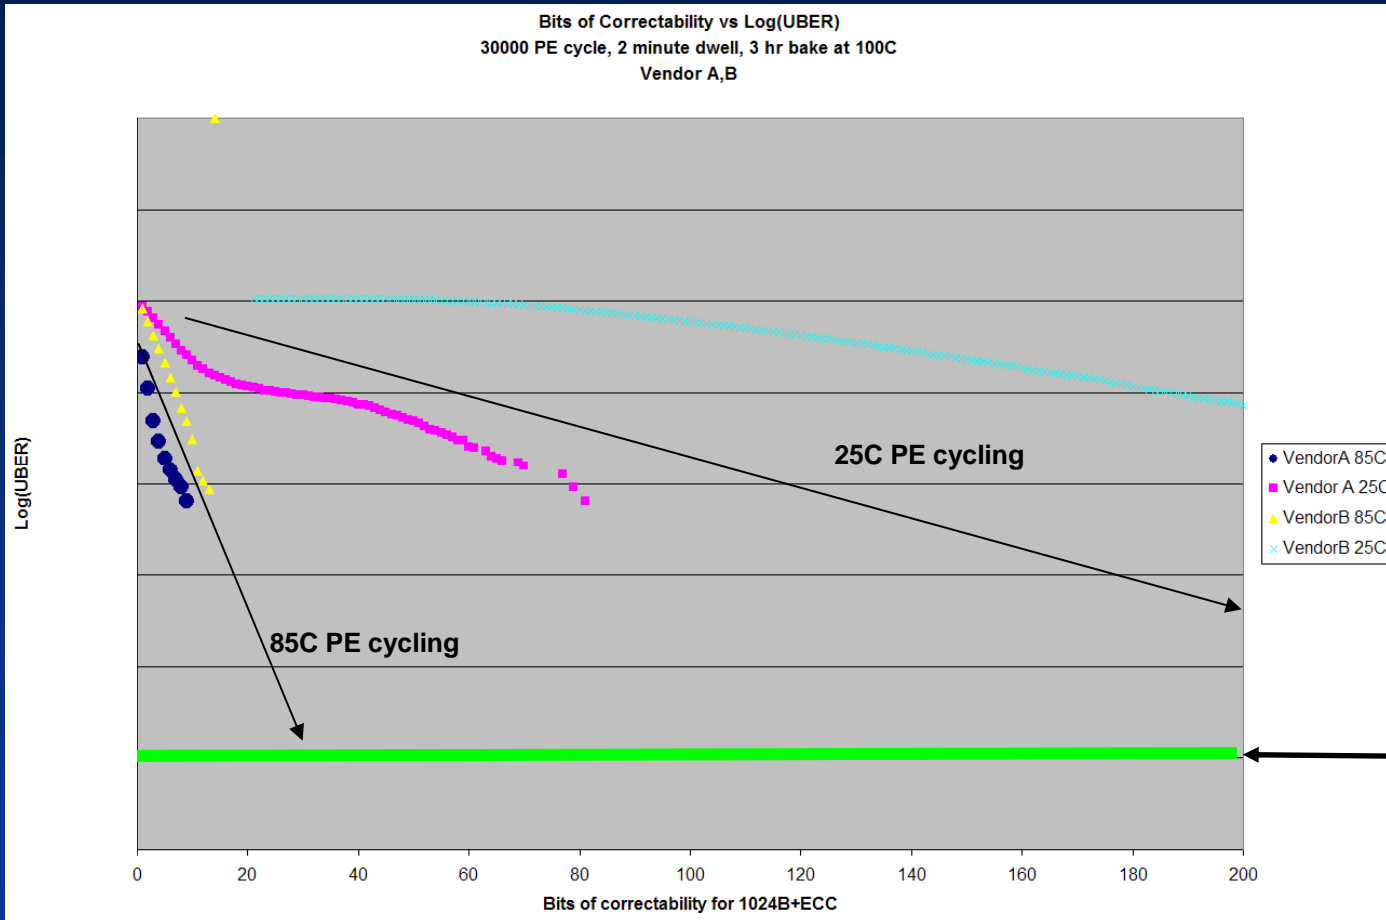
- As bake duration increases, raw bit error rate (RBER) increases
- Higher program-erase cycling temperatures results in improved data retention
- Longer dwell times show lower bit error rates

Post Program-Erase Cycling Data Retention Study - Dwell Time Sensitivity



- As dwell time increases, raw bit error rate (RBER) decreases
- Linear relationship on log-log plot implies a power law relationship – $RBER \sim DT^m$
 - Where m is a function of supplier, temperature, program-erase cycling and data retention bake duration
- Can be applied to predict RBER for extended dwell times (as in SSD environment)

Uncorrectable Bit Error Rate (UBER) Study

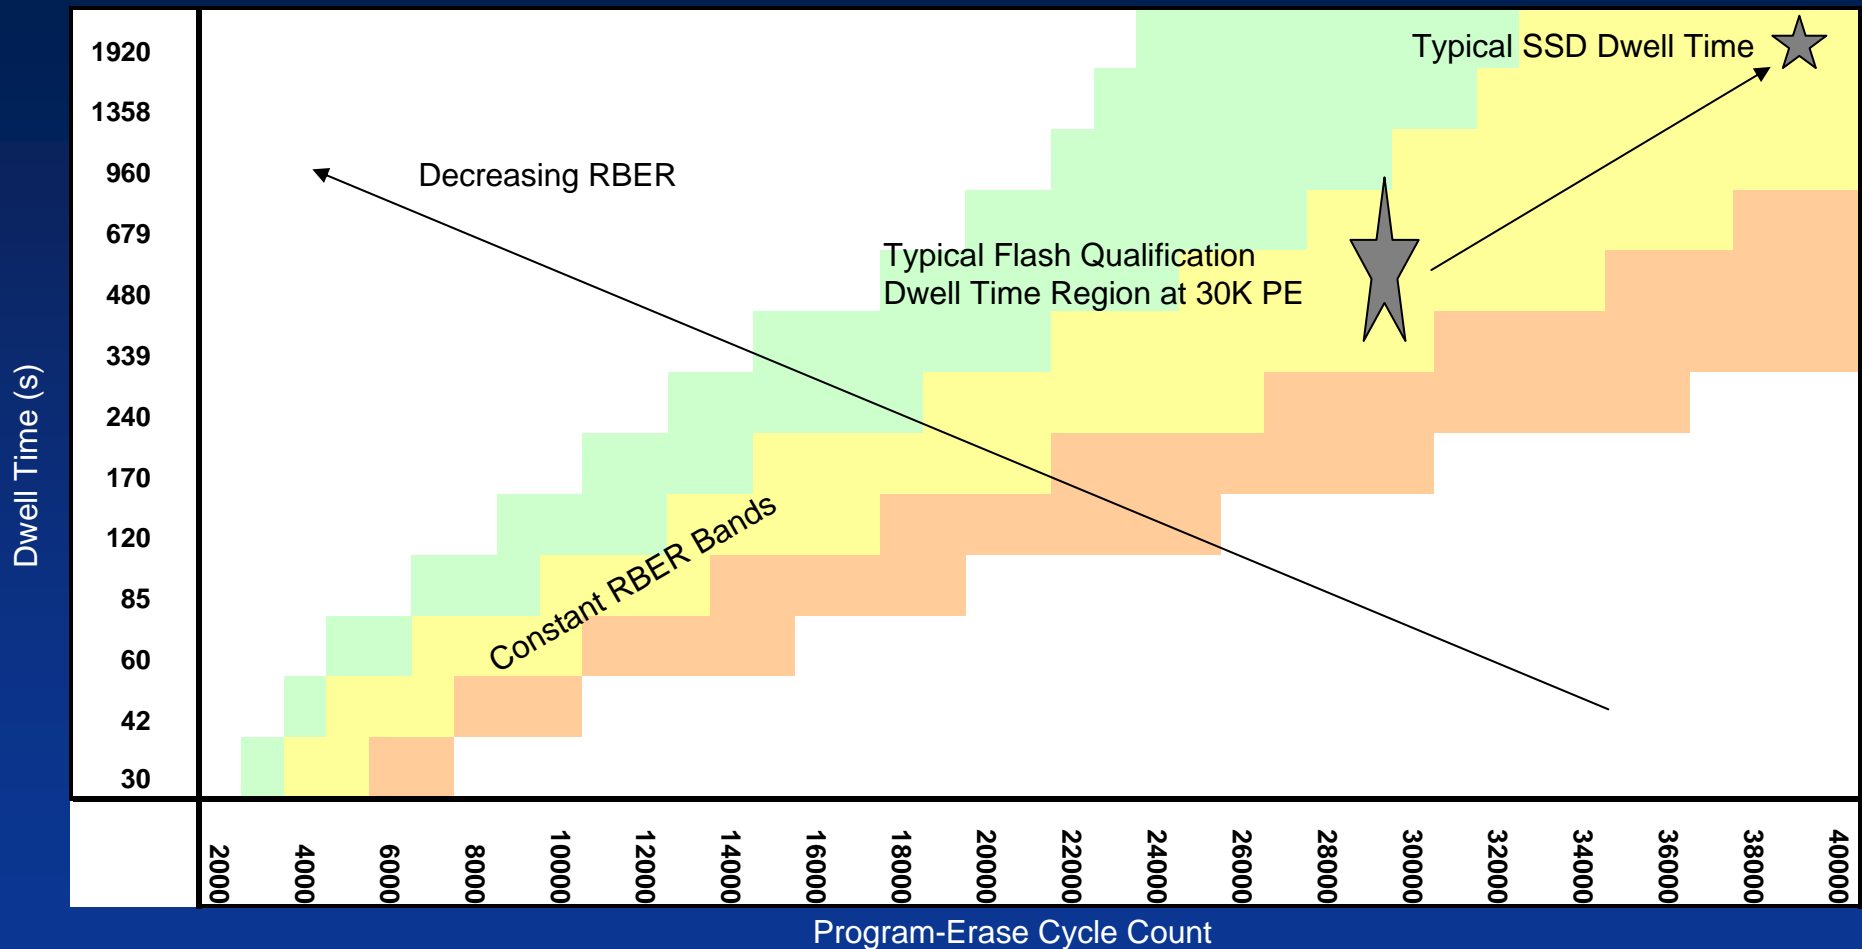


UBER = Number of 1KB UE
=====
total bits read

JEDEC SPEC
1E-14 (Fails / bits read)

- Data retention bit error rate is very sensitive to program-erase cycling temperature
- Devices cycled at higher temperatures have a lower RBER and lower ECC requirement
- For 30K program-erase cycled devices, 25C requires greater than 200 ECC bits per 1KB, while 85C requires about 35 ECC bits for the same data set

Extending Flash Dwell Time at Constant Data Retention



- Additional program-erase cycles can be realized when typical SSD-level dwell times are applied

Summary

- Flash device characterization pursued to describe relationship between variables and raw bit error rate, and to investigate Flash impact on SSD usable life
- Program-erase cycling shows some sensitivity to dwell time
 - Dependence observed for one supplier at 85°C (not seen at 25°C)
- Program-erase cycling shows sensitivity to temperature – suppliers under test show opposite relationships
- Post program-erase cycling data retention bit error rate shows clear sensitivity to dwell time during cycling
 - Extended dwell times exhibit lower bit error rate
- Post program-erase cycling data retention bit error rate shows clear sensitivity to temperature during cycling
 - Higher temperatures exhibit lower bit error rate
- Flash raw bit error rate vs. program-erase cycle dwell time log-log plot exhibits power law relationship
 - Can be applied to predict RBER for extended dwell times (as in SSD environment)
- Controller ECC requirement is reduced for higher program-erase cycling temperatures
- Additional Flash program-erase cycles can be realized when typical SSD-level dwell times are applied at constant data retention