

The Write Endurance of SSDs – Is MLC NAND Fit for Data Centers?

Moshe Selfin, Vice President Enterprise Solutions
Anobit

Storage Developer Conference
September 21st, 2010

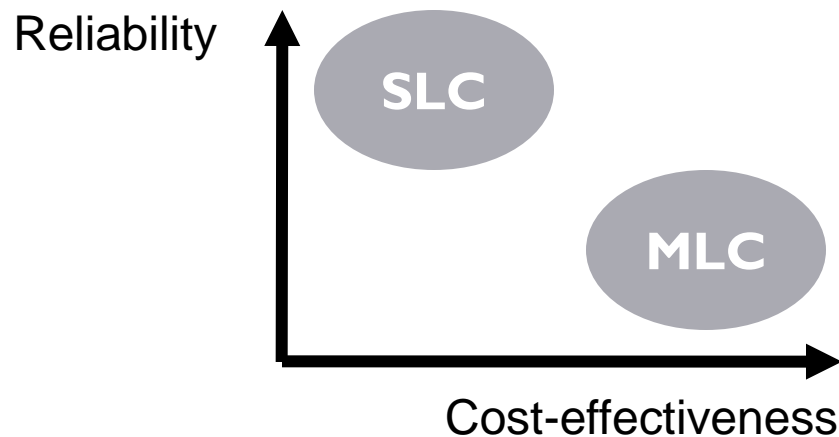
Agenda

- ❑ MLC vs. SLC NAND
- ❑ What does endurance really mean?
- ❑ How much endurance is really needed?
- ❑ How to make MLC fit for enterprise?

MLC vs. SLC NAND

	SLC	MLC	Factor
\$/GB*	\$6.75	\$1.41	X4.8
Endurance	50K-100K	3K	X20+
Failure Rate	100 PPM	5000PPM	X50

*32Gb Spot Price, Sep. 6th, 2010, www.memoryexchange.com



What Does Endurance Really Mean?

- **Flash Endurance:**
 - The number of times each flash cell can be programmed and erased before it is unusable
 - Specified in Program/Erase cycles
- **Drive Endurance:**
 - The number of times the full media can be written, before it is unusable
 - Can be specified in media cycles/day for 5 year


How to Translate Flash Endurance to Drive Endurance

Write Amplification

Amount of data written to the drive

Amount of data written by the user

- Write Amplification:
- Why Write Amplification?
 - The main reason: garbage collection
 - Efficient garbage collection can reduce write amplification by half



Over Provisioning

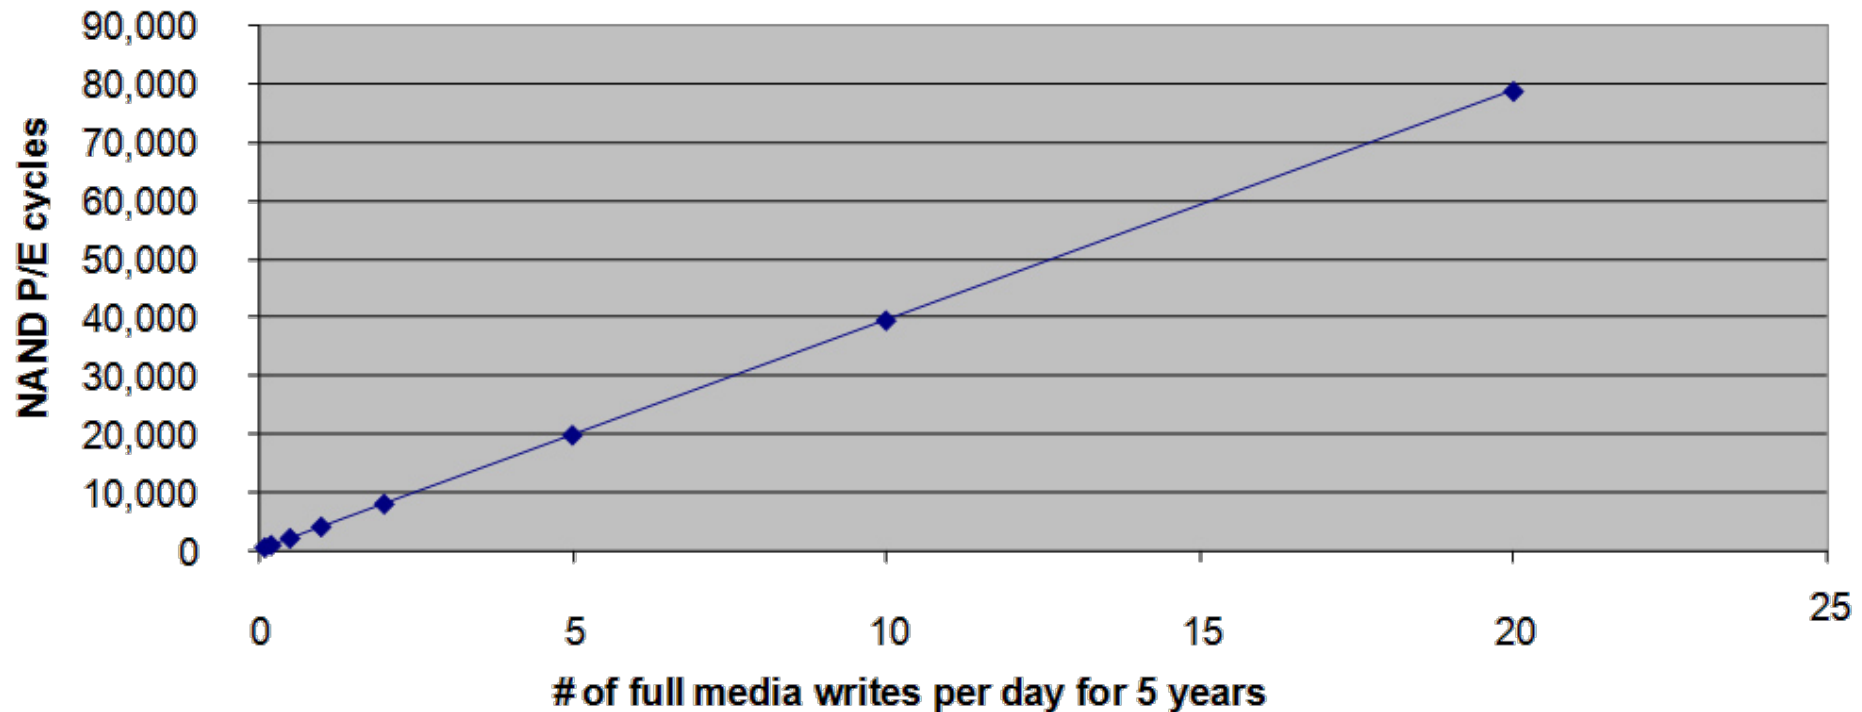
- Over Provisioning:
$$\frac{\text{Spare flash size}}{\text{Usable media size}}$$
- High over provisioning →
 - Larger average number of punches per block →
 - Less pages to copy →
 - Lower write amplification (higher efficiency)

Methods To Reduce Write Amplification

- ❑ Increase Over Provisioning
- ❑ Compression
 - ❑ The Benefit
 - ❑ Less data to write
 - ❑ The gained capacity is used for overprovision
 - ❑ The Downside - Data Dependent Pattern
 - ❑ Will not work when data is encrypted at application level
 - ❑ Will not work when most of the data is compressed (e.g. MS Windows)
- ❑ TRIM Command
 - ❑ In order to improve the overprovisioning, OS is requested to notify the drive which LBA are not used
 - ❑ Very basic TRIM support is implemented in Win 7, eMMC4.4 & Server 2008 R2

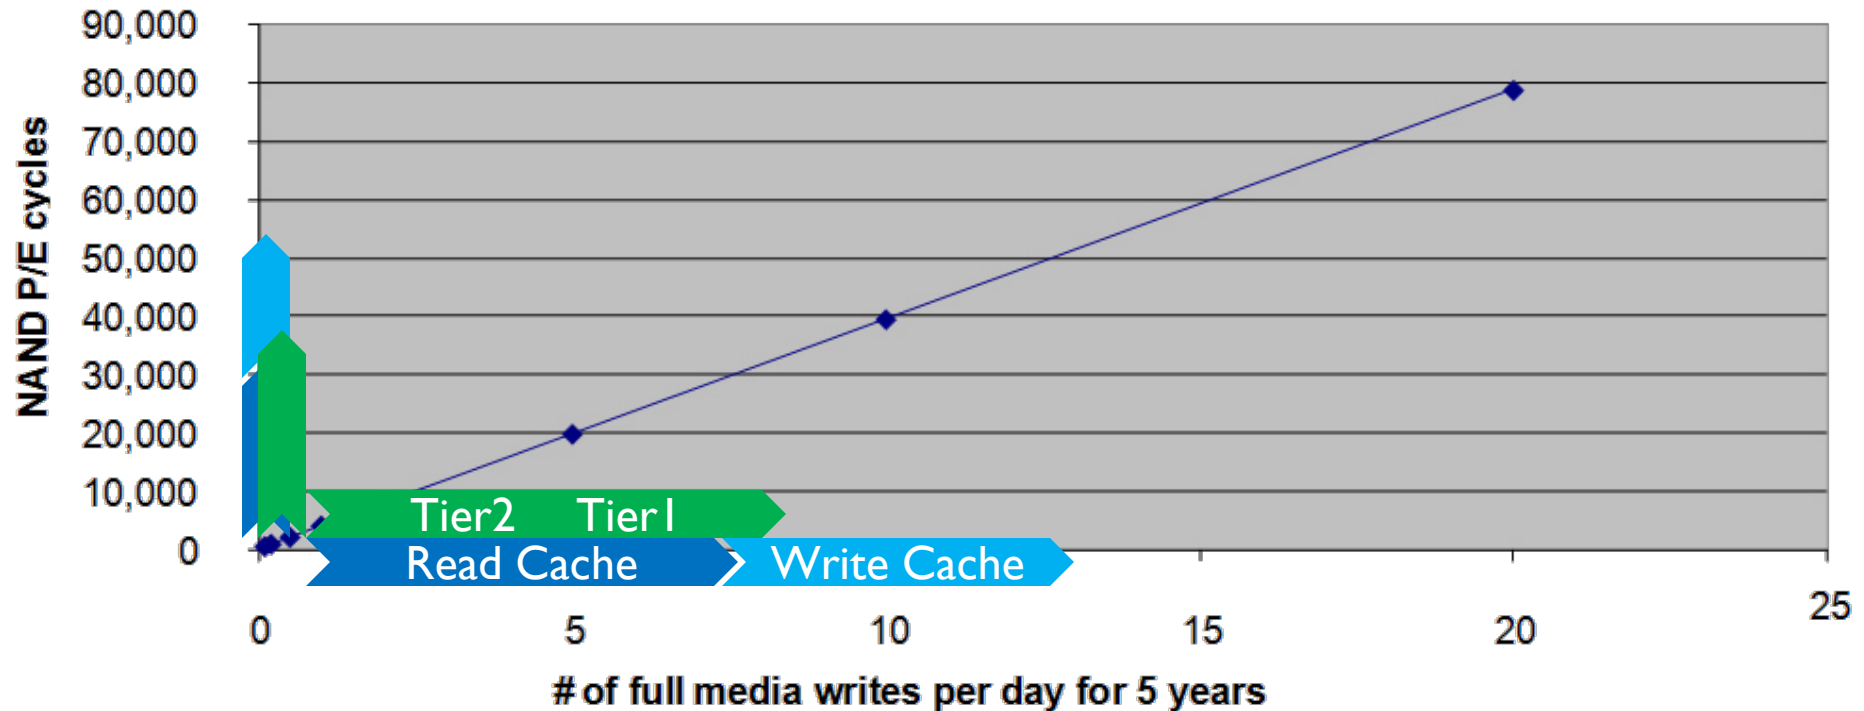
Drive Endurance vs. NAND Endurance

- Calculated for fully random, non-compressible data
- Assuming efficient garbage collection



How Much Endurance Is Really Needed?

Matching NAND and Application:




How to Make MLC Fit for Enterprise

"Enterprise MLC"

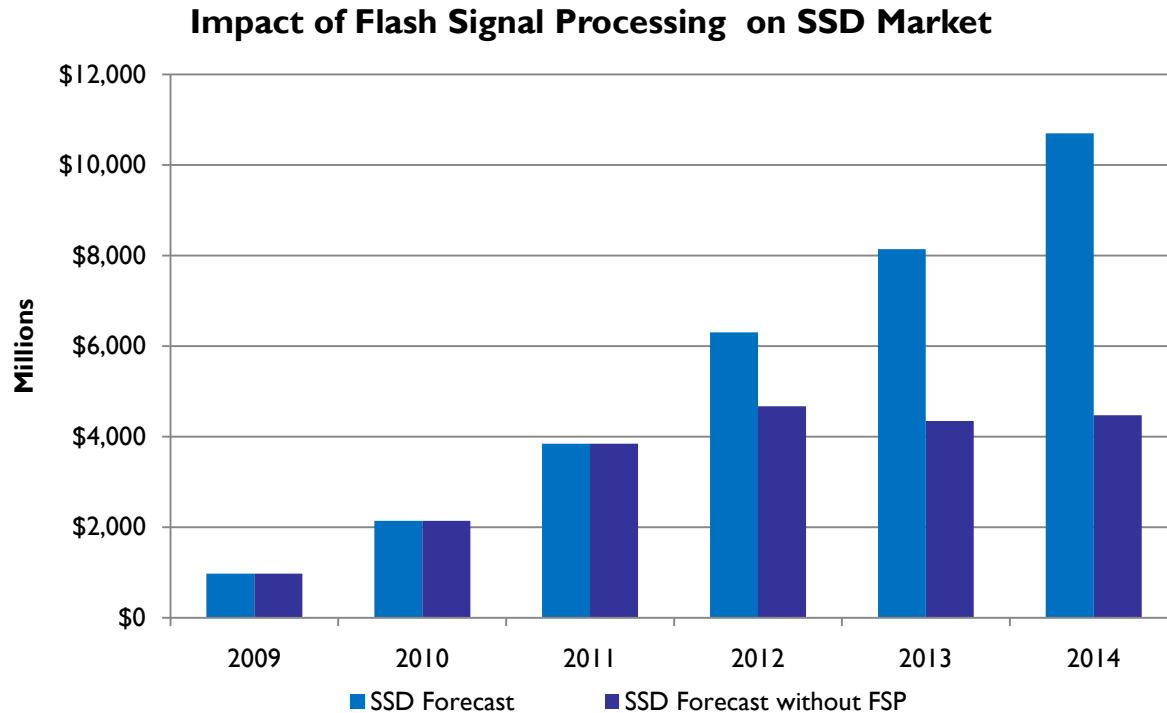
- ❑ High endurance MLC process
- ❑ Methods:
 - ❑ Screening parts, based on in-process variations
 - ❑ Reduced program performance
- ❑ Downside:
 - ❑ Requires special process handling
 - ❑ Enterprise SSD market is forecasted at 6% of the NAND market (~\$2B out of ~\$30B in 2015)
 - ❑ Will the NAND vendors commit a special process for this market size?
 - ❑ Reduced program performance

MLC + MSP – The "Holy Grail"

- ❑ Endurance: ~50K
- ❑ Reliability: ~100PPM
- ❑ No special process handling: **use consumer MLC**



Impact on the Industry



*Source: Forward Insights, August 2010

About Anobit

- Founded in 2006
- Based in Israel
- Subsidiaries in the US and Korea
- 130 Employees

Thank You