



Flash Memory Summit

# How Modern File Systems Detect & Respond to SD, SSD and other Managed NAND Media Failures

Thom Denholm  
Datalight



# Conflicting facts of life





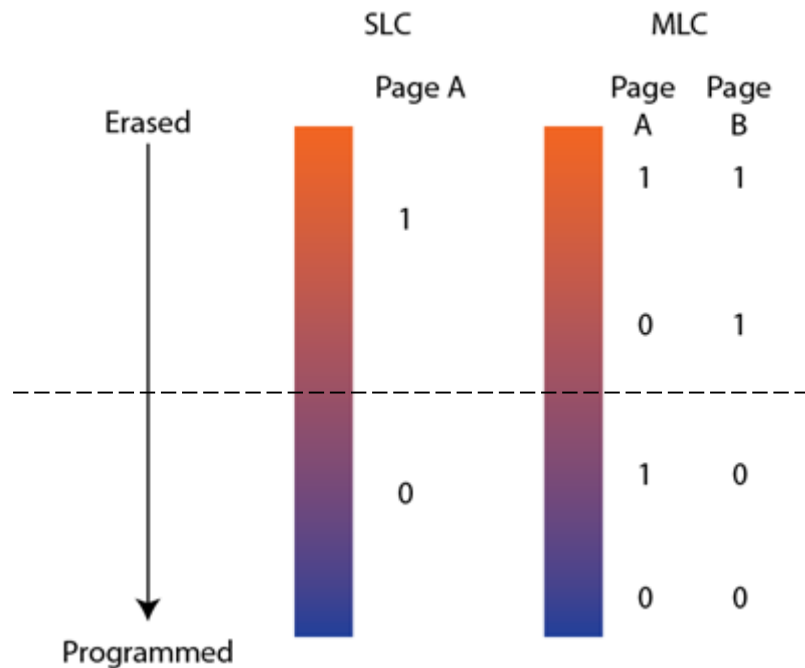
# Agenda

- NAND flash media and failures
- How software/firmware correct errors  
– but occasionally generate more!
- File system techniques to detect errors
- A special managed NAND error



# Types of NAND flash media


- Parallel / Serial (SPI)
- SLC / MLC (TLC)
- Ext./Int. Controller
- Firmware / Software





# Internal NAND Errors

- Write errors
- Bit errors
  - Error detection & correction (EDC)
- Power Interruption errors
  - Paired pages on MLC!





# Error Handling

## Hardware

- Maintain power during NAND write

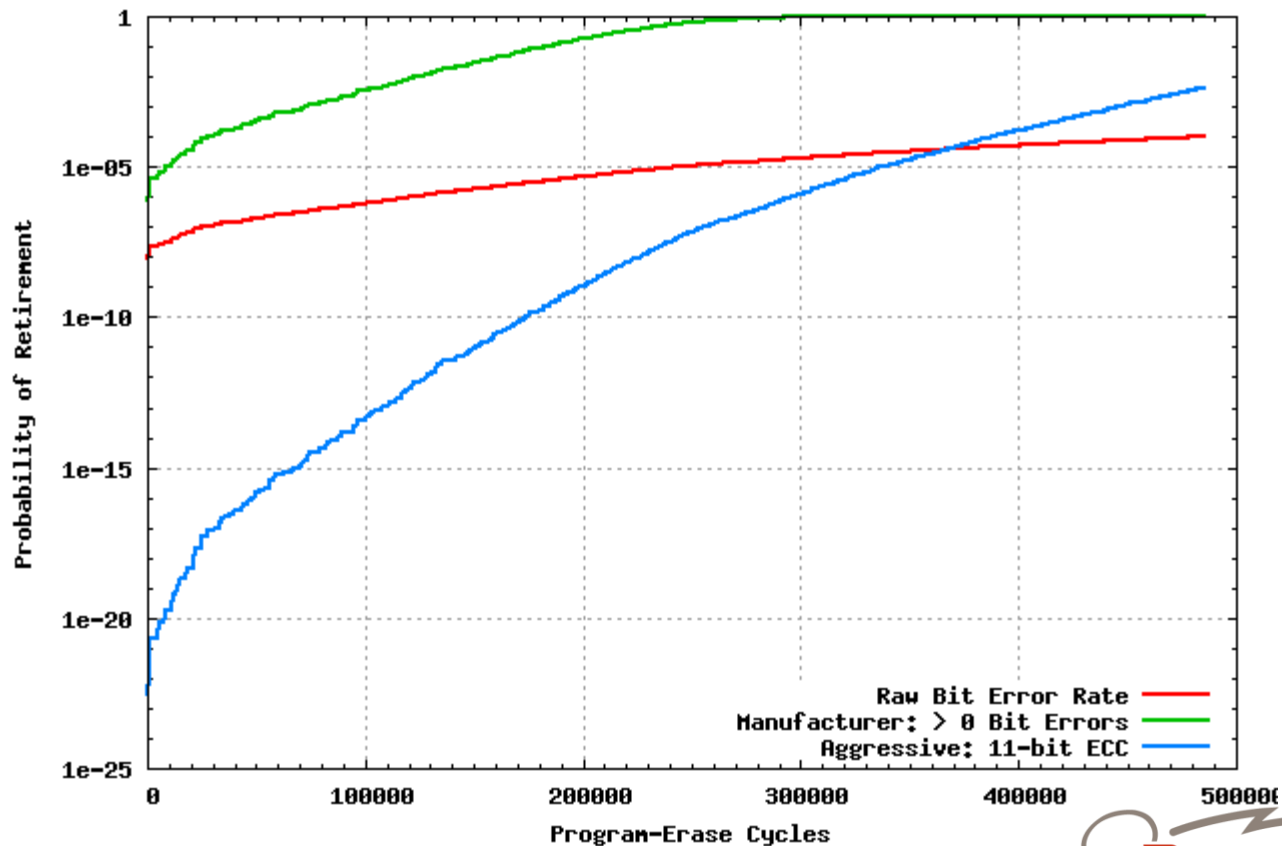
## Software

- Scrubbing of bit errors

- ~~Write Errors~~
- Bit Errors
- Power Interruption Errors



### ECC Comparison: Manufacturer Recommendation vs Aggressive Probability of PEB Retirement Vendor A, MLC





# File Systems

- General media access – nothing flash related
  - Exception – Discard interface
  
- Storage device failures
  - Write
  - Read





# Detecting media failure – CRC32

## File system metadata

- File names
- Dates & Times
- Location of file data
- CRC to validate contents





# CRC comparison chart


Linux FS	CRCs on		Release	Notes
	Metadata	File Data		
ext4	Some	No	3.5	Not on all metadata; not enabled by default
VFAT	No	No		
f2fs	No	No		
xf	Yes	No	3.2	
ZFS	Yes	Yes	3.14	ZFS uses the fletcher4 checksum instead of CRC32
btrfs	Yes	Yes	2.6.29	
Reliance Nitro	All	Optional	2.6.28	



## New situation - stale data

During customer image analysis

- Metadata CRCs matched
- All bits valid & correct
- Not part of the current file system state



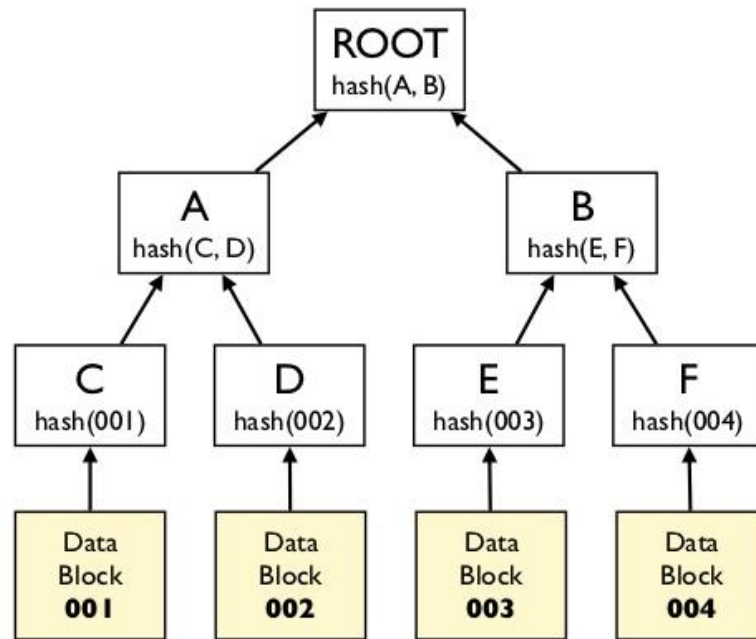
You can imagine the problems!



# Merkle Trees

## Hierarchical CRCs

- Both CRC32 and child node CRC
- Validates metadata is correct AND current





# Summary

- NAND flash media has errors
- Software/firmware can correct most of these
- File system must handle the rest!



Flash Memory Summit

# Questions?

Thom Denholm  
thom.denholm@datalight.com

Tuxera  
Booth #854

<https://www.datalight.com/resources/whitepapers/modern-file-systems-respond-to-media-failures>