

# DevAgent: Analyzing the Reliability of Storage Software Stack via Smart Devices

Mai Zheng, Computer Science Department, New Mexico State University zheng@nmsu.edu

### **Background & Motivation**

user

New NVM-based components are revolutionizing the traditional storage systems, potentially creating new failure modes difficult to understand

app/data	photos syste	m N	IVMKV library
	ext4 XFS	F2FS Rec	onFS DFS
OS kernel	single-queue block IO	multi-queue block IO	TRIM support NVMe driver
device	HDD	SSD buffer	
	flash	memory chins	sl

family payroll

• Failure example 1: 👩 algolia

June 2015, flash-based servers in Algolia data center started corrupting files; developers "spent a big portion of two weeks just isolating machines and restoring them as quickly as possible": Samsung SSDs were mistakenly blamed & blacklisted SAMSUNG (until one month later they identified a kernel bug)

healthcare records

• Failure example 2: SSDs exhibit different # of errors when tested on different OSes

OS (Kernel)	SSD 1	SSD 2	SSD 3
Debian 6 (2.6.32)	317	991	2
Ubuntu 14 (3.16)	88	0	1

testing higher layers

- Limitation of exiting analysis tools:
  - heavily rely on kernel & assume kernel is correct

### **Observations & Approach**

- We can no longer completely trust the kernel => analysis tools should minimize dependency/interference on kernel
- We can no longer focus on single component => cross-layer analysis
- We need to focus on generic & fundamental operations => interfaces b/w layers



## (minimal intrusion into kernel)

Identifying device-level impact of

kernel bug patches [3]

Prototyping DevAgent-FW on

Cosmos OpenSSD Platform [3]

- fd = open("/dev/sdb",...); user-level semantics fsync(fd); device-level 0x2A:WRITE semantics 0x35:SYNCHRONIZE CACHE
- Exposing vulnerabilities in popular file system checkers [1][2]



• Emulating realistic device behaviors for testing higher software layers [5][6]



 Analyzing fine-grained I/O behavior of parallel file systems [4]



(5) interrupted

(6) **fsck** 

- Publications:
  - [1] Understanding the Fault Resilience of File System Checkers (HotStorage'17)

DevAgent

**Results & Milestones** 

- [2] On Fault Resilience of File System Checkers (FAST'17- WiP)
- [3] Do Not Blame Devices for All Failures (NVMW'17-Poster)
- [4] A Generic Framework for Testing Parallel File Systems (PDSW'16)
- [5] Reliability Analysis of SSDs under Power Fault (TOCS'16)
- [6] Emulating Realistic Flash Device Errors with High Fidelity (NAS'16- Poster)

#### Acknowledgement

This material is based upon work supported in part by the National Science Foundation (NSF) under Grant Number 1566554 (CRII). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.



(7) repaired image