

IOWA STATE UNIVERSITY

Data Storage Lab



usenix  
ASSOCIATION

HotStorage '20

JULY 13-14, 2020


# On Failure Diagnosis of the Storage Stack

Duo Zhang, Om Rameshwar Gatla, Runzhou Han, Mai Zheng


Iowa State University



# Storage System Failures Are Troublesome



Oh No,  
I Lost My Data



# Existing Efforts Are Not Enough

- Mostly focus on *testing*
  - Require a special *testing environment*
    - e.g., a customized kernel
  - Still cannot prevent all failures in *production environment*

## Finding Semantic Bugs in File Systems with an Extensible Fuzzing Framework

Seulhee Kim, Meng Xu, Sanidhya Kashyap, Jungyeon Yoon, Wen Xu, Taesoo Kim  
Georgia Institute of Technology


**Abstract**  
File systems are too large to be bug-free. Although hand-written test suites have been widely used to stress the systems, they can hardly keep up with the rapid increase in the system size and complexity. Finding test cases being introduced and repeated regularly. These bugs occur in various flavors: simple buffer overflows to sophisticated semantic bugs. Although bug-specific checkers exist, they generally fail to explore the system states thoroughly. Here, we introduce a framework that enables the checking of various aspects of a file system under one umbrella, and highlights the potential of applying fuzzing to file systems but, in theory, any type of the system.

**1 Introduction**  
Designing and maintaining file systems are complicated. With the constant development for performance optimizations and new features, popular file systems have grown too large to be bug-free. For example, ext4 [1] and extfs [6], with 38k and 138k lines of code, respectively, witnessed 54 [2] and 113 [5] bugs reported in 2018 alone. A bug in a file system can wreak havoc on the user, as it not only corrupts data, but also poses security threats [3, 7, 8]. Therefore, it is crucial to explore the system states thoroughly. Here, we introduce a framework that enables the checking of various aspects of a file system under one umbrella, and highlights the potential of applying fuzzing to file systems but, in theory, any type of the system.

## Specifying and Checking File System Crash-Consistency Models

James Bornholt, Antonin Kaufmann, Balin Li, Arvind Krishnamurthy, Emin Torkal, Xi Wang  
University of Washington  
[bornholt, antonin.kaufmann, balinli, arvindk, emint, xiwang]@cs.washington.edu

**Abstract**  
Applications depend on persistent storage to recover state after system crashes. But the POSIX file-system interface does not define the possible outcomes of a crash. As a result, it is difficult for application writers to correctly understand the meaning of and dependencies between the system operations, which can lead to corrupt application state and, in the worst case, catastrophic data loss.



## EXPLODE: a Lightweight, General System for Finding Serious Storage System Errors

Jinfeng Yang, Cun Sun, and Dawson Engler  
Computer Systems Laboratory  
Stanford University

**Abstract**  
Storage systems such as the systems, databases, and RAID systems have a simple, basic contract: you give them data, they do not lose or corrupt it. Often they store the only copy, making an inevitable loss almost always bad. Unfortunately, their code is exceptionally hard to get right, since it must correctly manage data from any crash at any program point, on storage hardware that may be unreliable and/or present errors. This paper describes EXPLODE, a system that makes a systematically check and manage systems for errors. It provides a framework, generally specific specific checks and error messages, that can be used to check and manage systems. EXPLODE uses a novel algorithm to generate a comprehensive set of checks for storage systems.

**Abstract**  
dynamic storage checkers. EXPLODE makes it easy for checkers to find bugs in crash recovery code in their run on a live system; they tell EXPLODE when to generate the disk images that could occur if the system crashed at the current recovery point, which they then check for errors. We explicitly designed EXPLODE so that clients can check complex storage stacks built from many different subsystems. For example, Figure 1 shows a virtual test system on top of NFS on top of the ZFS filesystem on top of RAID. EXPLODE makes it easy for checkers for such deep stacks by providing a framework that lets users write small, checkable modules that plug them together to build a checker.

## Cross-checking Semantic Correctness: The Case of Finding File System Bugs

Changwoo Min, Sanidhya Kashyap, Byoungyeon Lee, Chengyu Song, Taesoo Kim  
Georgia Institute of Technology

**Abstract**  
Today, systems software is too complex to be bug-free. To find bugs in systems software, developers often rely on code checkers, like Linux's Sparse. However, the capability of existing tools used in commodity, large-scale systems is limited to finding only shallow bugs that tend to be introduced by simple programmer mistakes, and so do not require a deep understanding of code to find them. Unfortunately, the "deep" as well as those that are difficult to find are which violate high-level rules or invariants of the system. Thus, it is difficult for existing tools to find such bugs. This paper introduces a new checker, called Semantic Correctness, that can find such bugs. Semantic Correctness is implemented in C++ and is integrated with the Sparse checker.

## Finding Crash-Consistency Bugs with Bounded Black-Box Crash Testing


Jayashree Mohan<sup>1</sup>, Athile Martinez<sup>1</sup>, Soujanya Ponnampalli<sup>1</sup>, Pandian Raju<sup>1</sup>, Vijay Chidambaram<sup>1,2</sup>  
<sup>1</sup>University of Texas at Austin <sup>2</sup>VMware Research

**Abstract**  
We present a new approach to testing file-system crash-consistency: bounded black-box crash testing (B<sup>2</sup>CT). B<sup>2</sup>CT tests the file-system in a black-box manner using workloads of file-system operations. Since the space of possible workloads is infinite, B<sup>2</sup>CT bounds this space based on resources such as the number of file-system operations which operations to include, and exhaustively workloads within this bounded space. Each workload on the target file-system is simulated while the workload is being executed. The file-system receives a configurable number of concurrent test requests, upon receipt of which it must respond to the test requests. The test requests are generated by a test harness that simulates the behavior of a user process. The test harness is implemented in C++ and is integrated with the Sparse checker.




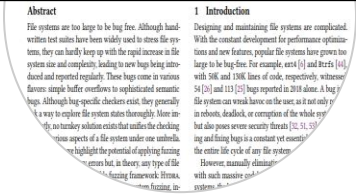
fact that even professionally managed databases occasionally suffer from power losses [19–22, 40, 41], it is important to ensure that file systems are crash-consistent. Unfortunately, there is little to no crash-consistency testing today for widely-used Linux file systems such as ext4, xfs [25], btrfs [21], and ZFS [23]. The current practice in the Linux file-system community is to use of any proactive crash-consistency testing. If a user reports a crash-consistency bug, the file-system developers then reactively write a test to capture that bug. System developers use `xfsctest` [1] to test for correctness tests, some of which are implemented in C++ and are integrated with the Sparse checker.

# Existing Efforts Are Not Enough

- Mostly focus on *testing*
  - Require a special *testing environment*
    - e.g., a customized kernel
  - Still cannot prevent all failures in *production environment*

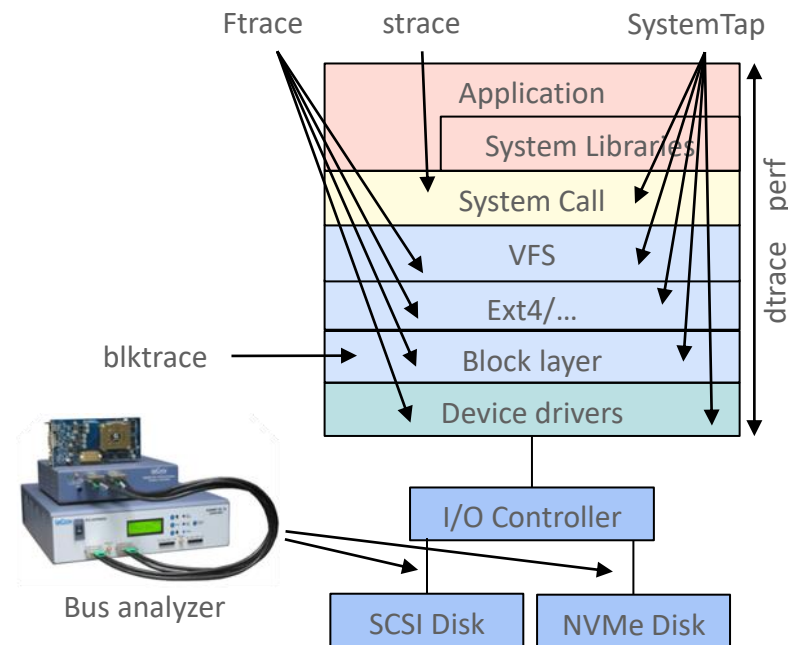


# What to do if failures happen?



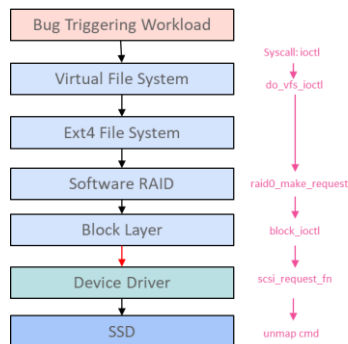
# Practical Diagnosis Tools & Limitations

- Practical diagnosis tools
  - Software-based
    - e.g., GDB, SystemTap, Ftrace
  - Hardware-based
    - e.g., Bus analyzer
- Limitations
  - Require substantial manual efforts
    - e.g., GDB single-stepping
  - Require special hardware
  - Only cover partial storage stack



# A Real-World Case: Diagnosis Is Challenging

- Algolia data center incident:
  - Servers **crashed** and files **corrupted** for unknown reason
  - After weeks of diagnosis, Samsung SSDs were **mistakenly blamed**
  - After one month**, a Linux kernel bug was identified as root cause



## When Solid State Drives are not that solid

Adam Surak | Jun 15th 2015 | 12 min read | Engineering

[BLOG](#) [Algolia.com](#) [Product overview](#) [Latest releases](#) [Resources](#)

f

The level of despair was reaching a critical level and the pages in the middle of the night were unstoppable. We spent a big portion of two weeks just isolating machines as quickly as possible and restoring them as quickly as possible. The one thing we did was to implement a check in our software that looked for empty blocks in the index files, even when they were not used, and alerted us in advance.

in

[BLOG](#) [Algolia.com](#) [Product overview](#) [Latest releases](#) [Resources](#)

f

As a result, we informed our server provider about the affected SSDs and they informed the manufacturer. Our new deployments were switched to different SSD drives and we don't recommend anyone to use any SSD that is anyhow mentioned in a bad way by the Linux kernel. Also be careful, even when you don't enable the TRIM explicitly, at least since Ubuntu 14.04 the explicit `FSTRIM` runs in a cron once per week on all partitions – the freeze of your storage for a couple of seconds will be your smallest problem.

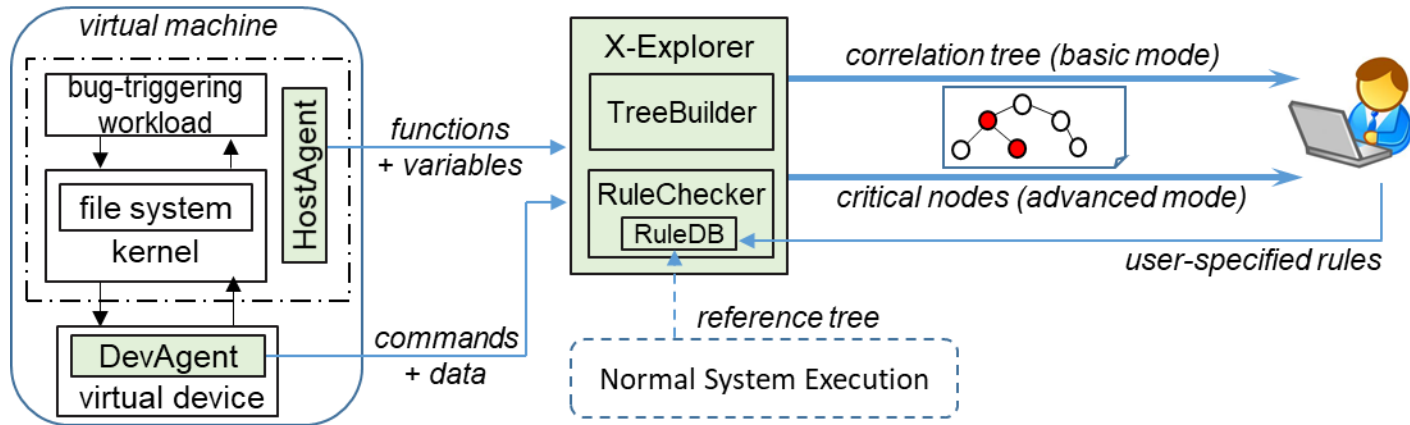
in

t

# Our Approach

# X-Ray: A Cross-Layer Approach

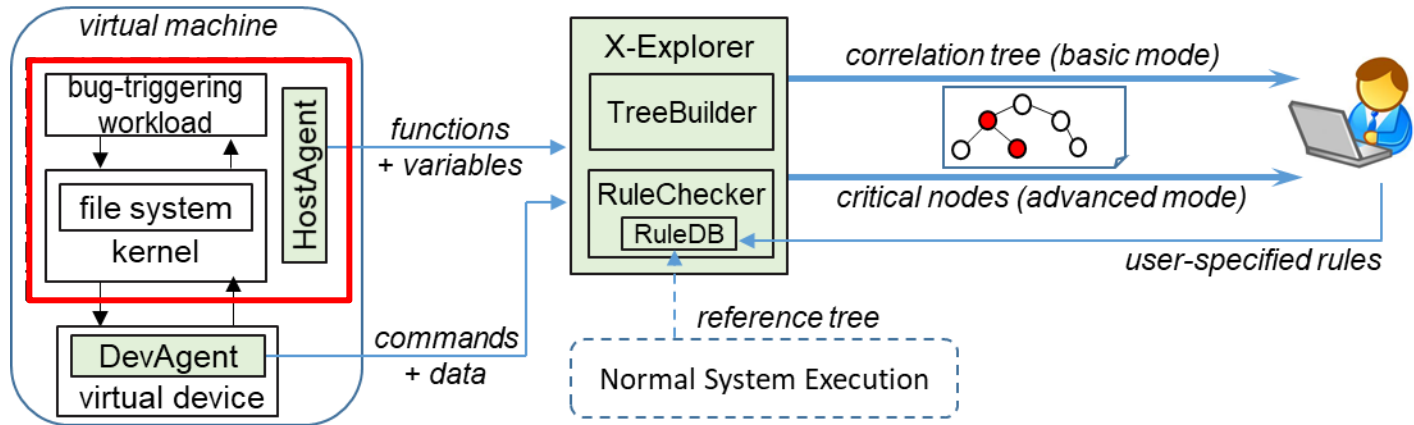
- Support unmodified software stack
- Intercept device activity without relying on kernel or special hardware
- Visualize multi-layer correlation
- Narrow down root cause (semi)automatically





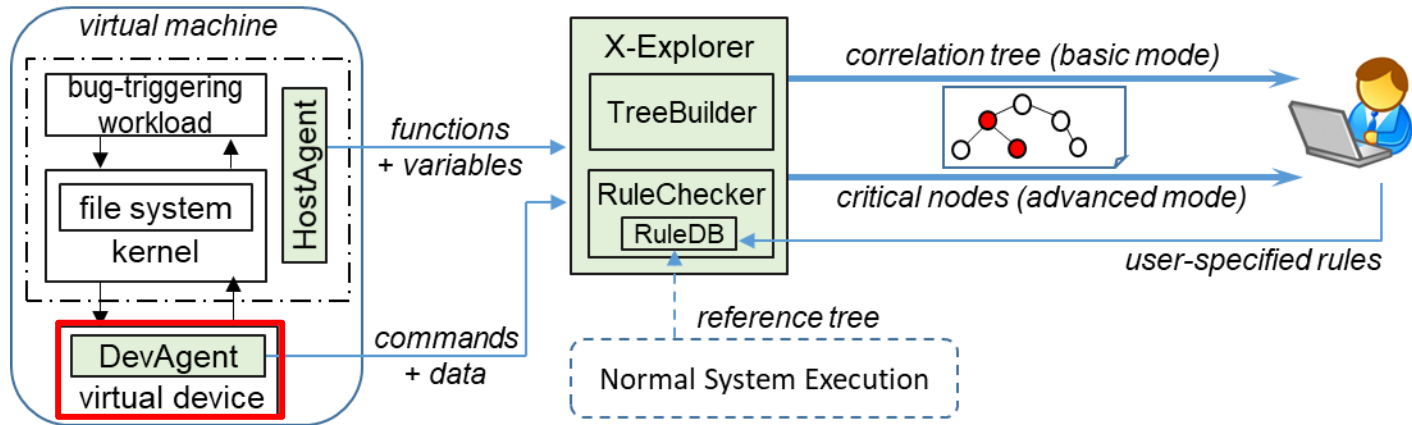
# X-Ray: A Cross-Layer Approach

- HostAgent: help understand host-side system activities
  - Trace host-side events
    - e.g., syscalls, kernel functions



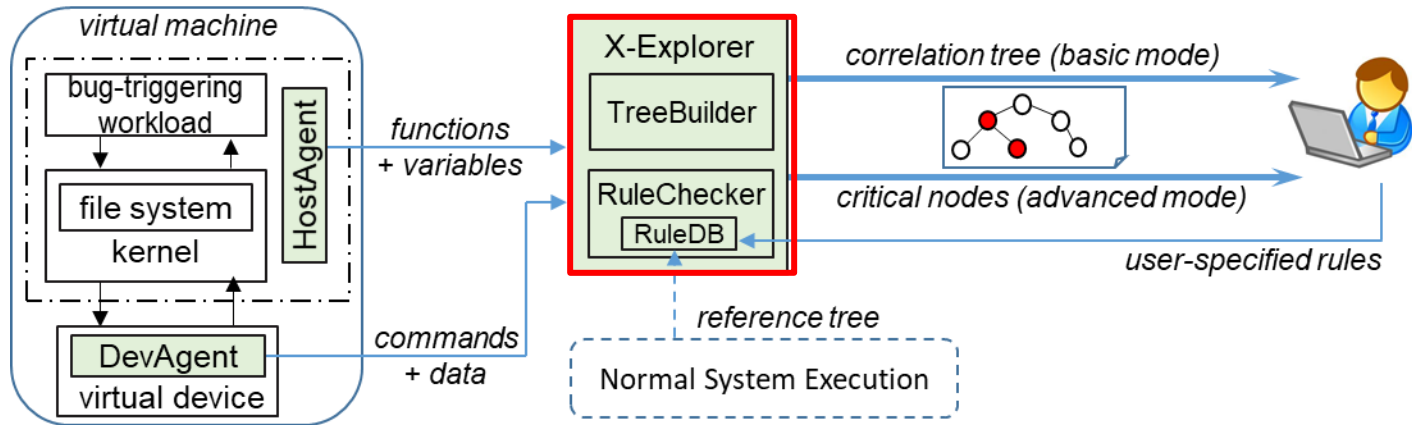
# X-Ray: A Cross-Layer Approach

- DevAgent: help understand changes of persistent states
  - Trace device commands
    - e.g., SCSI, NVMe




# X-Ray: A Cross-Layer Approach

- X-Explorer: facilitate diagnosis in two ways
  - Build and visualize multi-layer correlation (i.e., correlation tree)
  - Highlight critical nodes/paths based on rules




# Key Challenge #1

- How to correlate information across layers ?



# Key Challenge #1

- How to correlate information across layers ?
  - Cannot use SCSI/NVMe hints ☹️
    - Require modification to workload/OS
  - Use timestamp 😊
    - Customized Ftrace frontend
      - Convert execution time to epoch time
    - NTP(Network Time Protocol) based synchronization
      - Solve accuracy problem caused by virtualization



# Key Challenge #2

- How to reduce manual efforts ?



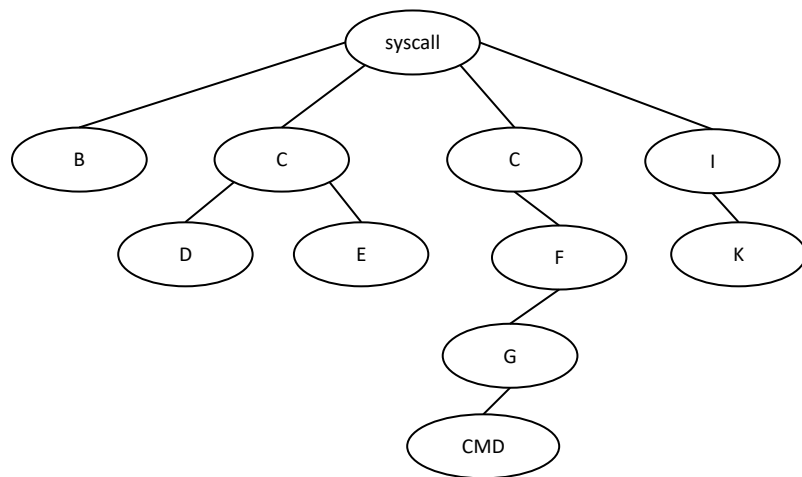
# Key Challenge #2

- How to reduce manual efforts ?
  - Visualize cross-layer events & dependencies in a correlation tree

Dependency

Syscall	→	B
Syscall	→	C
C	→	D
C	→	E
Syscall	→	C
C	→	F
F	→	G
G	→	CMD
Syscall	→	I
I	→	K

Tracing log



Cross-layer tree




# Key Challenge #2

- How to reduce manual efforts ?
  - Visualize cross-layer events & dependencies in a correlation tree
  - Automatically narrow down the root cause via rules



# Key Challenge #2


- How to reduce manual efforts ?
  - Visualize cross-layer events & dependencies in a correlation tree
  - Automatically narrow down the root cause via rules
    - Rules specified by users (e.g., “ancestors of device commands”)



Correlation tree



Rule specified by users →





Critical part





# Key Challenge #2

- How to reduce manual efforts ?
  - Visualize cross-layer events & dependencies in a correlation tree
  - Automatically narrow down the root cause via rules
    - Rules specified by users (e.g., “ancestors of device commands”)
    - Rules derived from reference execution (e.g., non-failure run due to different kernel version)



Rules derived from reference



Tree from Abnormal execution    Tree from reference execution




Difference part



# Preliminary Results

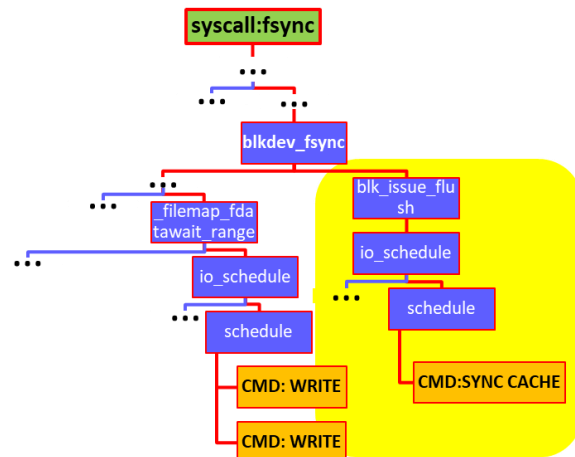
# Preliminary Results

- Case Study
  - A kernel bug manifested as serialization errors on SSDs [Zheng *et. al.*@TOCS'16, FAST'13]
  - The problem can be observed in the correlation tree clearly
  - Rules can help narrow down the root cause quickly



Tree from abnormal execution

Rules →



Pinpointed root cause

# Preliminary Results

- Result summary
  - 5 failure cases reported in the literature
  - 3 simple rules to define critical parts of the correlation trees
  - Reduce the search space for root causes effectively
    - 0.06% - 4.97% nodes of the original trees

Case ID	node count in original tree	node count by Rule#1	node count by Rule#2	node count by Rule#3
1	11,353 (100%)	704 (6.20%)	571 (5.03%)	30 (0.26%)
2	34,083 (100%)	697 (2.05%)	328 (0.96%)	22 (0.06%)
3	24,355 (100%)	1254 (5.15%)	1210 (4.97%)	/
4	273,653 (100%)	10230 (3.74%)	/	/
5	284,618 (100%)	5621 (1.97%)	5549 (1.95%)	/

# Conclusion and Ongoing Work

- X-Ray: A cross-layer approach for failure diagnosis
  - Support unmodified software stack
  - Intercept device activity without relying on kernel or special hardware
  - Visualize multi-layer correlation
  - Narrow down root cause (semi)automatically
- Explore more real-world failure cases
- Derive more diagnosis rules
- Automate the comparison based on reference tree

# Thanks !

Duo Zhang

[duozhang@iastate.edu](mailto:duozhang@iastate.edu)

<https://www.ece.iastate.edu/~mai/lab/dsl.html>

