


# Evaluating File System Reliability on Solid State Drives

Shehbaz Jaffer, **Stathis Maneas**, Andy Hwang, Bianca Schroeder



USENIX ATC '19 Lightning Talk

# Introduction & Motivation

- Storage landscape has changed:

# Introduction & Motivation

- Storage landscape has changed:
  - ~~HDDs~~ -> SSDs.

# Introduction & Motivation

- Storage landscape has changed:
  - ~~HDDs~~ -> SSDs.
  - Different failure characteristics:

# Introduction & Motivation

- Storage landscape has changed:
  - ~~HDDs~~ -> SSDs.
  - Different failure characteristics:
    - Partial failures are a magnitude higher for SSDs!

# Introduction & Motivation

- Storage landscape has changed:
  - ~~HDDs~~ -> SSDs.
  - Different failure characteristics:
    - Partial failures are a magnitude higher for SSDs!
  - New/Changed file systems:

# Introduction & Motivation

- Storage landscape has changed:
  - ~~HDDs~~ -> SSDs.
  - Different failure characteristics:
    - Partial failures are a magnitude higher for SSDs!
  - New/Changed file systems:
    - ext3 -> **ext4** (journaling).
    - **Btrfs** (copy-on-write).
    - **F2FS** (log-structured, tailored for flash).

# Introduction & Motivation

- Storage landscape has changed:
  - ~~HDDs~~ -> SSDs.
  - Different failure characteristics:
    - Partial failures are a magnitude higher for SSDs!
  - New/Changed file systems:
    - ext3 -> **ext4** (journaling).
    - **Btrfs** (copy-on-write).
    - **F2FS** (log-structured, tailored for flash).
- How do these file systems deal with partial drive errors?



# Research Questions & Methodology

# Research Questions & Methodology

## **What we know** (IRON File Systems, 2005):

- Only ext3, JFS, ReiserFS, NTFS (partial).
- Hard disks only.
- Does not consider file system checkers.

# Research Questions & Methodology

## What we know (IRON File Systems, 2005):

- Only ext3, JFS, ReiserFS, NTFS (partial).
- Hard disks only.
- Does not consider file system checkers.

## What we *want* to know:

- Btrfs, ext4, F2FS.
- Can they **detect** errors?
- Can they **recover** from errors?
- Can their **system checker fix** errors?

# Research Questions & Methodology

## What we know (IRON File Systems, 2005):

- Only ext3, JFS, ReiserFS, NTFS (partial).
- Hard disks only.
- Does not consider file system checkers.

## What we *want* to know:

- Btrfs, ext4, F2FS.
- Can they **detect** errors?
- Can they **recover** from errors?
- Can their **system checker fix** errors?

# Research Questions & Methodology

## What we know (IRON File Systems, 2005):

- Only ext3, JFS, ReiserFS, NTFS (partial).
- Hard disks only.
- Does not consider file system checkers.

## What we *want* to know:

- Btrfs, ext4, F2FS.
- Can they **detect** errors?
- Can they **recover** from errors?
- Can their **system checker fix** errors?

# Research Questions & Methodology

## What we know (IRON File Systems, 2005):

- Only ext3, JFS, ReiserFS, NTFS (partial).
- Hard disks only.
- Does not consider file system checkers.

## What we *want* to know:

- Btrfs, ext4, F2FS.
- Can they **detect** errors?
- Can they **recover** from errors?
- Can their **system checker fix** errors?

## Our Approach

# Research Questions & Methodology

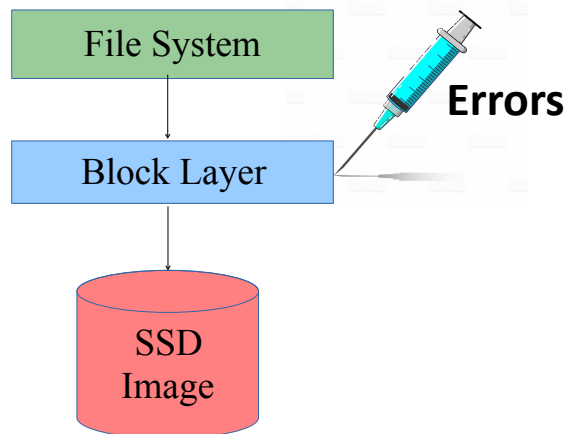
## What we know (IRON File Systems, 2005):

- Only ext3, JFS, ReiserFS, NTFS (partial).
- Hard disks only.
- Does not consider file system checkers.

## What we *want* to know:

- Btrfs, ext4, F2FS.
- Can they **detect** errors?
- Can they **recover** from errors?
- Can their **system checker fix** errors?

## Our Approach



# Research Questions & Methodology



## What we know (IRON File Systems, 2005):

- Only ext3, JFS, ReiserFS, NTFS (partial).
- Hard disks only.
- Does not consider file system checkers.


## What we *want* to know:

- Btrfs, ext4, F2FS.
- Can they **detect** errors?
- Can they **recover** from errors?
- Can their **system checker fix** errors?


## Our Approach




Test Programs



Read/Write/  
Corruption  
Experiments



Data Structures



Recovery &  
Detection



**7,128 Test Cases**





# Results Overview

File System	Detection	Recovery	Observations





# Results Overview

File System	Detection	Recovery	Observations
ext4			





# Results Overview

File System	Detection	Recovery	Observations
ext4			Data loss due to shorn and lost writes.







# Results Overview

File System	Detection	Recovery	Observations
<b>ext4</b>			Data loss due to shorn and lost writes.
<b>Btrfs</b>			







# Results Overview

File System	Detection	Recovery	Observations
<b>ext4</b>			Data loss due to shorn and lost writes.
<b>Btrfs</b>			Several cases lead to kernel crashes or panics.

# Results Overview

File System	Detection	Recovery	Observations
<b>ext4</b>			Data loss due to shorn and lost writes.
<b>Btrfs</b>			Several cases lead to kernel crashes or panics.
<b>F2FS</b>			

# Results Overview

File System	Detection	Recovery	Observations
<b>ext4</b>			Data loss due to shorn and lost writes.
<b>Btrfs</b>			Several cases lead to kernel crashes or panics.
<b>F2FS</b>			<ul style="list-style-type: none"><li>• Does not detect and report write errors.</li><li>• Can become unmountable due to data corruption.</li><li>• System checker has room for improvement.</li></ul>