

# The Security War in File Systems: An Empirical Study from A Vulnerability-Centric Perspective

JINGHAN SUN, University of Illinois Urbana-Champaign, USA

SHAOBO LI, University of Illinois Urbana-Champaign, USA

JUN XU, University of Utah, USA

JIAN HUANG, University of Illinois Urbana-Champaign, USA

This paper presents a systematic study on the security of modern file systems, following a vulnerability-centric perspective. Specifically, we collected 377 file system vulnerabilities committed to the CVE database in the past 20 years. We characterize them from four dimensions that include why the vulnerabilities appear, how the vulnerabilities can be exploited, what consequences can arise, and how the vulnerabilities are fixed. This way, we build a deep understanding of the attack surfaces faced by file systems, the threats imposed by the attack surfaces, and the good and bad practices in mitigating the attacks in file systems. We envision that our study will bring insights towards the future development of file systems, the enhancement of file system security, and the relevant vulnerability mitigating solutions.

## 1 INTRODUCTION

After decades of development since the 1960s, file systems have become a core component in nearly all computer systems. Working as the direct interfaces to user data and typically running at a high privilege, file systems are highly security-sensitive. This has incentivized the creation of various techniques to harden the security of file systems, ranging from formal verification [20, 22, 71, 85] to access control [29, 73] and sanity checking [33, 40, 52].

Despite the above efforts, the security of file systems is still frequently compromised [9, 72], causing damages such as data breaches and hijacked execution. To gain insights towards changing this situation, it is beneficial to build a systematic understanding of the security of file systems, from angles such as the attack surfaces in file systems and the threats imposed by the attack surfaces. Past studies [47, 60] on file systems mostly focus on functionality-related perspectives, failing to provide the desired security-centric understanding. Prior study [15, 47] took an initial effort to investigate the filesystem bugs and security in Linux. However, it targeted solely on generic file systems, with a focus on the causes and consequences, which lacks the analysis of the whole lifecycle of the vulnerability exploitation and fixes processes. And we believe a thorough understanding of the attack exploitation procedure in file systems could help developers to advance the design principles and implementation strategies of file systems.

In this paper, we present an in-depth study on the security of various file systems, aiming to fundamentally understand the entire vulnerability exploitation procedure and possible fix strategies. Specifically, our study centers around addressing the following questions:

- *Q1: What are the major attack surfaces in file systems?*
- *Q2: What are the threats posed by the attack surfaces?*
- *Q3: What are the common approaches to exploiting these attack surfaces by attackers?*
- *Q4: How are the attack surfaces fixed?*

**Methodology:** To answer the above questions, our study takes a vulnerability-centric approach. We collected 377 file system vulnerabilities committed to the CVE database [31] in the past 20 years. They reported explicit security implications and covered nearly all the mainstream file systems (*generic file systems* like Ext4 [50], XFS [77], and F2FS [42], *mobile file systems* like HFS [36] and APFS [11], and *networked file systems* like NFS [56], GlusterFS [27], and Ceph [18]).

We manually analyzed the 377 vulnerabilities and sought answers to questions *Q1 - Q4*. First, we investigated the root causes of the vulnerabilities by analyzing the vulnerable code and the

corresponding patch(es). We then categorized the vulnerabilities based on their root causes. This enabled us to gain an understanding of the attack surfaces opened by the filesystem vulnerabilities and the characteristics of the attack surfaces (*answering Q1*). Second, we manually and statically reasoned the possible execution paths to exploit the vulnerabilities. In this process, we summarized the potential threats that adversaries can bring by following those exploiting paths (*answering Q2 and Q3*). Finally, we surveyed the patches to the vulnerabilities, unveiling the good and bad practices in mitigating the attack surfaces (*answering Q4*). Throughout the study, we derived a set of findings. We summarise the major ones that help answer *Q1-Q4* as follows.

**Attack surfaces:** Vulnerabilities in file systems are diverse in types. Although 27% of the vulnerabilities are of general kinds that may appear in other software systems (e.g., memory errors and race conditions), a unique cause of these general vulnerabilities is the addition of new features, such as crash-consistency models into file systems. The remaining vulnerabilities are unique to file systems but attributed to different reasons:

- (1) 29% are rooted in sanity checks that are used to validate the states of various file system properties (e.g., namespace, inode attributes, and file category). The sanity checks are highly semantic related, whose implementation is complicated and often leads to vulnerabilities.
- (2) 15% are due to the lack of permission checks or incorrect access control. Similar to sanity checks, permission checks and access control are semantic dependent and highly complex. Their buggy implementations largely turn into permission and access-related vulnerabilities.
- (3) 18% are introduced when the functionality is extended to support the use environments of specific file systems. For instance, many vulnerabilities appear in encryption mechanisms adopted by mobile file systems to defend against physical attacks.

**Threats and their exploitation procedure:** File systems vulnerabilities are dangerous. As we will illustrate in Figure 6, file system vulnerabilities pose a similar security threat level to that of vulnerabilities in other OS components. The specific consequences after exploiting file system vulnerabilities vary across vulnerability types.

- (1) Exploiting general types of vulnerabilities (e.g., memory errors) in file systems can largely cause their common sequences, such as denial-of-service (DoS) and hijacked execution. An interesting observation is that general vulnerabilities in file systems can more often lead to data leakage because file systems directly deal with user data.
- (2) Exploiting erroneous sanity checks can also lead to DoS but more often data leakage, which allows attackers to initiate further attacks such as privilege escalation and permission bypass.
- (3) Exploiting lack of permission checks or incorrect access control produces more predictable consequences, often including permission bypass and privilege escalation as well as subsequent damages, such as data leakage and data corruption.
- (4) Exploiting the remaining vulnerabilities typically causes consequences pertaining to the use contexts of the file systems. For instance, exploiting vulnerabilities in the encryption modules of mobile file systems often enables physical attacks, while exploiting vulnerabilities in the networking management of network file systems usually leads to code injection attacks.

**Fixes with patches:** File system patches are diverse in both complexity and functionality. 56% of the patches only change less than 10 lines of code, using simple operations such as adding a boundary check and fixing a syntax error (see the details in Figure 8). In contrast, around 20% of the patches bring modifications to over 50 lines of code, involving complex behaviors ranging from full validation of data structures to fixes in multiple components. The majority of existing and current patches work in an ad-hoc style. For instance, the patches to many race conditions simply make the responsible operations atomic. These ad-hoc patches can heal the target vulnerabilities, but they fail to prevent similar vulnerabilities from emerging or fundamentally mitigate them. Nevertheless,

Table 1. File systems covered in our study.

Type	File Systems	# of CVE
Generic	Ext4, VFS, XFS, F2FS, BtrFS, profs, Ext2, Ext3, tmpfs, ReiserFS, JFS	178
Mobile	HFS/HFS+, APFS, F2FS, eCryptFS	57
Networked	NFS, OpenAFS, GlusterFS, CephFs, HDFS	142
<b>Total</b>	-	377

our study identifies positive signs. Strategies to enforce preventive and secure-by-design patching are arising. For instance, many patches of race conditions switch from lock-based implementations to lock-free implementations, fundamentally eliminating issues like deadlocks. File systems are also often correlated, so as their patches. In particular, a patch to the VFS subsystem often requires updating the underlying file systems. Otherwise, the patch may not come into effect. However, today’s practice often overlooks the correlated file systems when patching the target ones, thus leading to insecure patches.

**Contributions:** Our main contributions are as follows.

- We perform the first in-depth study on the security of file systems, from a vulnerability-centric perspective. The study unveils the attack surfaces of file systems, followed by demystifying the root causes and threats of the attack surfaces.
- We present 21 previously-unknown findings throughout the study. Besides facilitating a deeper understanding of the attack surfaces, the findings also reveal how the development of file systems influences or even impedes security.
- We distill 7 insights from the results and findings of the study. The insights point out the deficiencies in today’s practice of addressing security threats faced by file systems and also bring up a set of guidance towards improvements.

## 2 STUDY METHODOLOGY

Our study aims to understand the attack surfaces in file systems, from the perspectives summarized as *Q1 - Q3*. To support the study, we consider public vulnerabilities as the data source. In comparison to other applicable data sources such as attacks that happened in the past, vulnerabilities are more accessible and more abundant.

### 2.1 Data Collection

We collected file system vulnerabilities from the CVE database [31]. We considered the CVE database for three reasons. First, the CVE reports are reviewed and confirmed by trustworthy parties (major IT vendors, security companies, and research organizations [7]) and thus, offer high reliability. Second, the CVE reports largely provide important details, such as root cause information and patch information. Such details are helpful for us to answer *Q1 - Q3*. Third, the CVE reports are large in number and diverse in terms of sources, which should properly reflect the overall situation.

Specifically, we examined the CVE reports committed in the past 20 years (January 1999 to December 2020) and identified those associated with mainstream file systems. Older reports were omitted since they may not represent what is happening at present. We further refined the reports to only keep those carrying information about or giving references to (i) location of the vulnerable code, (ii) analysis of the root cause (comprehensive or brief), (iii) possible exploitation of the vulnerability and subsequent consequences (which will be verified following our approach in §2.2), and (iv) the fixes with patches.




Fig. 1. Distribution of vulnerabilities over the past twenty years.

In the end, we collected 377 file system vulnerabilities, whose distribution across time is shown in Figure 1. They cover most of the widely deployed file systems, as illustrated in Table 1. The vulnerabilities also span a wide spectrum of types and related details are presented in §3.

## 2.2 Vulnerability Analysis

We ran three analyses on each of the 377 vulnerabilities.

**Root cause analysis.** The first analysis focuses on understanding the root cause of a CVE report. In this analysis, we first extract descriptions of the root causes from the report and then verify the description through interpretation of the vulnerable code and its patch.

**Exploitation analysis.** Following the root cause analysis, we aim to understand the possible exploitation of a vulnerability and the subsequent consequences. One common approach is to reproduce the existing proof-of-concept (PoC) exploits against the vulnerability [54]. This approach, however, has limitations. First, PoC exploits are not available for many CVE reports. Omitting those CVE reports can downgrade the generality of our studies. Second, the reproduction of PoC exploits is labor intensive, which cannot scale to support large-scale studies like ours.

In our study, we use a more generic and more lightweight approach. We first parse the description in a CVE report to retrieve the reported exploitation and its consequences. We then reason the vulnerable code to identify execution paths that can verify the description. Specifically, an execution path successfully verifying the exploitation must meet the following conditions:

- The execution path covers a feasible route from a user-space interface to the vulnerable code.
- The execution path covers a feasible route to triggering the vulnerability.
- The execution path covers a feasible route from triggering the vulnerability to causing the described consequence(s).

If no such execution paths can be identified for a CVE report, we skip the report. The example in Figure 2 illustrates our approach. In the example, a local data structure *fa* (declared at line 2) is not fully initialized before it is copied to user space at line 7. The CVE report says “It allows local users to obtain potentially sensitive information from kernel stack memory via an ioctl call”.

We find that any execution path first opening a legitimate file in an XFS file system and then calling *ioctl* (with *XFS\_IOC\_FSGETXATTR*) on the opened file will trigger the vulnerability, and leak data from the kernel stack. This way the exploitation and the consequence are considered verified.

**Patch analysis.** The last analysis examines the patches to file system vulnerabilities. One goal is to understand the strategies used by the existing patches (e.g., ad-hoc or systematic). This way, we expect to unveil whether today’s practice of file system patching can help fundamentally resolve the issues of vulnerabilities. Another goal of the analysis is to summarise the aspects overlooked

---

```

1 int xfs_ioc_fsgetxattr(xfs_inode_t *ip,int attr,...){
2     struct fsxattr fa;
3     + memset(&fa, 0, sizeof(struct fsxattr));
4     // data structure `fa' is uninitialized before patching
5     xfs_ilock(ip, XFS_ILOCK_SHARED);
6     fa.fsx_xflags = xfs_ip2xflags(ip);
7     if (copy_to_user(arg, &fa, sizeof(fa)))
8         return -EFAULT;
9 }

```

---

Fig. 2. An uninitialized memory vulnerability in XFS and its patch (CVE-2010-3078).

by file system developers when patching vulnerabilities. We anticipate bringing related evidence to incentivize broader attention to those aspects.

### 2.3 Threats to Validity

While we carefully designed our methodology, it may still pose threats to the validity of our findings. First, similar to other sampling-based studies, our study can miss representative vulnerabilities or even miss certain types of vulnerabilities. Thus, our findings might be biased towards the collected samples. We admit the possibility of this threat but we should have reduced it by adopting the most comprehensive, most standard dataset. Second, we verify the exploitation of a vulnerability and the consequences through static reasoning. Unlike PoC exploits, static reasoning does not guarantee the fidelity of the results. As such, we may present underestimated/exaggerated/erroneous threats of the vulnerabilities. To mitigate this threat, we document the execution paths identified to exploit the vulnerability and then have another person review and vet the execution paths. Any execution paths that raise discrepancy are discarded. Third, the analysis of the patches can often require deep domain knowledge. Lack of such knowledge will lead to misunderstanding about the patches. To mitigate the threat, we only allow authors whose research focus is file/storage system to look at the patch and again, we have multiple authors to vet the results.

## 3 TAXONOMY OF FILE SYSTEM VULNERABILITIES

Vulnerabilities appear in every mainstream file system and their discovery also presents an increasing trend, as shown in Table 2 and Figure 1. To gain better insights, the first deeper understanding we build about is the categories of vulnerabilities. We categorize the vulnerabilities into two types:

- *Generic vulnerabilities* – vulnerabilities that can appear in any kinds of software systems. The major types include memory errors and race conditions.
- *Semantic vulnerabilities* – vulnerabilities that only appear in file systems but can appear in all kinds of file systems. These vulnerabilities are mostly caused by cursoriness or errors in semantic-dependent or functionality-dependent implementations.

The categorization brings a set of statistical findings and interesting observations.

**F-1: Generic vulnerabilities are common in file systems (>27%) but present positive signs. And there has been a misconception that generic vulnerabilities are decreasing as file systems become mature.** Our study reveals that there would be a significant amount of deep memory and concurrency vulnerabilities that can be detected with advanced bug-finding tools, as shown in Figure 3. Generic vulnerabilities account for a significant portion of the vulnerabilities in file systems. In particular, memory errors and concurrency issues led to 21.7% and 5.5% of all the vulnerabilities, respectively. The numbers are not surprising since file systems are largely developed in memory unsafe languages (e.g., C/C++) and widely support concurrency.

Table 2. Summary of file system vulnerabilities. *M*: Memory Errors; *C*: Concurrency Issues; *S*: Sanity Check Errors; *P*: Permission Errors; *N*: Network Errors; *O*: Others (e.g., hash collisions like CVE-2014-7283).

Type	FILE SYSTEMS				VULNERABILITIES					
	Name	Release Time	#CVEs	M	C	S	P	N	O	
Generic	JFS [2]	1990	4	50%	0%	50%	0%	0%	0%	
	Ext2 [17]	1993	7	14%	14%	43%	14%	0%	14%	
	XFS [77]	1993	23	30%	4%	35%	4%	0%	26%	
	VFS [4]	1995	28	18%	11%	40%	18%	0%	14%	
	procfs [53]	1999	18	22%	6%	11%	22%	0%	39%	
	ReiserFS [5]	2000	5	40%	0%	40%	20%	0%	0%	
	Ext3 [76]	2001	11	9%	0%	36%	18%	0%	36%	
	tmpfs [74]	2001	6	17%	0%	17%	33%	0%	33%	
	Ext4 [50]	2008	45	13%	4%	49%	4%	0%	24%	
	Btrfs [64]	2009	13	15%	8%	38%	31%	0%	8%	
	F2FS [42]	2012	18	11%	11%	56%	6%	0%	17%	
<b>Total</b>	<b>Avg →</b>	178	19%	6%	39%	13%	0%	22%		
Mobile	HFS/HFS+ [36]	1985	8	25%	0%	75%	0%	0%	0%	
	eCryptFS [34]	2006	16	44%	0%	25%	19%	0%	13%	
	F2FS [42]	2012	11	36%	0%	45%	9%	0%	9%	
	APFS [11]	2016	4	25%	25%	25%	25%	0%	0%	
	Others	-	18	22%	17%	33%	6%	0%	22%	
<b>Total</b>	<b>Avg →</b>	57	32%	7%	39%	11%	0%	12%		
Networked	NFS [59]	1984	72	17%	7%	13%	20%	24%	19%	
	GlusterFS [27]	2005	18	16%	0%	11%	11%	39%	22%	
	OpenAFS [3]	2006	27	44%	4%	11%	11%	22%	7%	
	HDFS [35]	2006	6	0%	0%	33%	67%	0%	0%	
	CephFS [79]	2012	19	16%	0%	16%	21%	26%	21%	
<b>Total</b>	<b>Avg →</b>	142	21%	4%	13%	19%	25%	16%		




Fig. 3. Distribution of memory errors reported from file systems over time. The deployment of JANUS [82] and Code Property Graph (CPG) [83] led to spikes of reported memory bugs in 2013 and 2018.

No doubt, the above numbers give the important message that generic vulnerabilities still constitute a major attack surface in file systems. Apparently, it is unwise to underestimate or overlook generic vulnerabilities at this point. However, we do observe positive signs. Various bug-finding tools (e.g., JANUS [82] and Code Property Graph [83]) and bug-detecting tools (e.g., AddressSanitizer [67] and ThreadSanitizer [68]) emerged in recent years. Real-world deployment




Fig. 4. Distribution of vulnerabilities across time. Generic vulnerabilities are stabilizing and semantic vulnerabilities present an increasing trend.

of these tools has been disclosing generic vulnerabilities in file systems. Consider Figure 3 as an example. Code Property Graph [83] and JANUS [82], after deployment, discover a surge of memory errors in generic file systems. JANUS applies fuzzing techniques to file system vulnerability detection by mutating file system images and manipulating file system operations as the input to the fuzzer. The large search space of fuzzing helps detect deep memory errors that cannot be manually tested via hand-crafted images (e.g., CVE-2018-10880). This could encourage researchers to properly apply other advanced bug detection tools to file systems.

**F-2: Semantic vulnerabilities are the dominant category of vulnerabilities in file systems (43%).** File systems involve highly complex semantics for reliable and efficient file management. Implementation of the semantics is challenging and often error-prone, bringing the dominant category of vulnerabilities. There are two major sources leading to semantic vulnerabilities, sanity checks and file permissions. We separately discuss them in the following.

*Sanity checks.* File systems maintain a mass of semantic-related states, such as file system namespace, inode attributes, cache consistency, and so on. Missing sanity checks on the states is a major way of introducing vulnerabilities. For instance, the implementations of the journaling in Ext2 and Ext3 never validate the journal superblock before reading it, which can cause an assertion or even kernel corruption (CVE-2011-4132). Erroneous sanity checks on the states is another way of introducing vulnerabilities. For example, the `f2fs` utilities before version 1.12.0 has a logical flaw when validating the sanity of the superblock (CVE-2020-6070). The logical flaw can lead to bypass of validation and even user-controlled execution. In total, missing sanity checks and erroneous sanity checks caused 29% of the total vulnerabilities in file systems.

*File permissions.* File permissions are indispensable to the protection of user data. To enforce file permissions, file systems rely on standard permission models (e.g., `u/g/o:r/w/x`) and access control list (ACL) (e.g., POSIX ACLs) to manage accesses. The actual implementations of permission models and ACL, however, often miss or add insufficient permission checks, leading to 15% of the vulnerabilities in file systems. For example, ReiserFS misses permission checks on the accesses to the directory storing file attributes (i.e., `xattrs`) (CVE-2010-1146). This enables any user to access or modify the attributes of files that are not allowed.

In contrast to generic vulnerabilities, semantic vulnerabilities are more concerning today. On the one hand, most of the semantic vulnerabilities seem to be discovered in an ad-hoc manner. We did not identify many tools that systematically detect semantic vulnerabilities. While some generic tools, like the fuzzing-based JANUS, do help find semantic vulnerabilities, they are still better at detecting generic vulnerabilities as they often lack domain-specific guidance to capture semantic vulnerabilities. On the other hand, as shown in Figure 4, the number of reported semantic vulnerabilities is clearly increasing over time. This might be attributed to that more file systems are included in our study as time moves forward or that the number of hidden semantic vulnerabilities

Table 3. Distribution of vulnerabilities caused by different features in block management. Only representative file systems are selected. The numbers are per file system.

File Systems	Extent	Extended Attribute	Delayed Allocation	Flex/Meta Block Group	Inline Data
Ext4 [50]	29%	13%	4%	16%	6%
XFS [77]	6%	16%	-	-	-
F2FS [42]	-	6%	-	-	-
JFS [2]	-	3%	-	-	-

is significantly large. But either way, the trend implies an increasing threat posed by semantic vulnerabilities.

**Summary:** Both generic vulnerabilities and semantic vulnerabilities open large attack surfaces in file systems. However, generic vulnerabilities present a seemingly improving trend while semantic vulnerabilities show the opposite.

**Insight-1:** To heal the attack surfaces of file systems, more efforts should be prioritized to mitigate semantic vulnerabilities. In particular, it is wise to explore tools that can systematically unveil semantic vulnerabilities.

Our categorization brings knowledge about the major vulnerabilities to be addressed. However, it gives few insights into how. To this end, we further investigate which file system functionality and what file operations introduce the vulnerabilities. In a nutshell, every major functionality of file systems brings in vulnerabilities, but inode management, block management, and page cache system brings in more. And a key observation is that extending the existing core functionality with new features contributed to a substantial number of vulnerabilities. We detail our findings below.

**F-3: Inode management is a major source of introducing vulnerabilities (13% out of Linux FS).** As per our study, inode management is also one of the most vulnerable parts of file systems. One reason is the complexity in the nature of inode management. Another reason, which is becoming more dominant, is the increasing need to support newly emerging features such as inline data blocks and extended attributes. A particular example is Ext4. Unlike Ext2 and Ext3 that uniformly keep extended attributes in a block, Ext4 supports three on-disk formats to store extended attributes, including internal inode, additional block, and dedicated inode. In regular cases where the extended attributes are small, Ext4 simply stores them in the internal inode. However, when the attributes grow bigger, Ext4 will swap them to an additional block. This type of format adjustment brings extra complexities and causes multiple vulnerabilities (e.g., CVE-2018-11412). Prior work proposed isolating metadata structures of unrelated files and directories to avoid failure propagation [48]. And inode replication was also proposed to ensure its fault tolerance [6]. We believe more efforts should be made to decouple the *inode* management from other components in file systems.

**F-4: Block management incurs plenty of vulnerabilities (21%) and presents an increasing trend. These vulnerabilities could be mitigated by having a more expressive interface to offload block management to the storage devices.** Block management is another major source of introducing vulnerabilities. In Table 3, we present the ratio of vulnerabilities caused by different features in block management. This is not surprising since block management has complexity comparable to inode management. Moreover, more features are being integrated into block management for better efficiency. In particular, recent file systems, such as Ext4 and BtrFS, use extents instead of direct/indirect blocks to organize file blocks. The codebase of extent management (7K lines of code) is much larger than that of direct/indirect block organization (1.5K lines of code). The need for fast indexing, flexible block group [23] and delayed block allocation all contribute to the increasing complexity of the block management, and hence result in more vulnerabilities.



Instead of purely relying on the storage software for the block management, we believe that future file systems should explore an expressive interface to offload the block management tasks to storage devices. This could not only ease the file system development and verification procedure, but also isolate its vulnerabilities from the OS kernel. As the hardware resources (i.e., storage processor and memory capacity) inside storage controllers become more powerful, the device is capable of handling more management tasks. The most recent work on device-level file systems, such as DevFS [39] and KEVIN [41], demonstrated the feasibility of offloading the file system functionalities to storage devices, however, they purely focused on performance improvement. It is highly desirable to explore the design trade-offs with respect to the security enhancement.

**F-5: Page cache causes 12% of the vulnerabilities and they could leak sensitive kernel information.** Modern file systems widely use page cache to temporarily store data blocks that will be flushed to the persistent storage at certain time points. The use of page cache avoids frequent interactions with the persistent storage and thus, improves the performance of applications. However, the extra layer of cache creates new space for vulnerabilities. A major type of such vulnerabilities is caused by uninitialized cache. Specifically, developers often do not initialize many fields of metadata structures saved in page cache. Later when the page cache is flushed, the uninitialized fields, whose memory may still carry sensitive kernel information, will be leaked to the persistent storage (e.g., CVE-2005-0400 and CVE-2006-6054). Any user who has access to the block can steal the sensitive information. As we will discuss in §4, uninitialized memory in cache systems poses higher threats than uninitialized memory in general, because the former may always propagate to the disk at the point of flush. Therefore, the memory safety for page cache is especially critical in file systems.

**F-6: Crash-consistency models for persistent storage often lead to uninitialized memory.** Many file systems use crash consistency models, such as journaling [76], logging [65], and shadow paging [37]. For instance, JFS employs a synchronous writing strategy to log the storage operations and inode. This will ensure that the file system can always recover to a correct state at a crash. The support of crash-consistency introduces new data structures and requires sophisticated synchronizations across different components, which involves high complexity and often incurs vulnerabilities. For example, JFS logs all relevant in-memory data structures but many of them may not be properly initialized. The uninitialized data may later be written to persistent storage, resulting in leakage of kernel information (e.g., CVE-2004-0181).

**F-7: The mismatch between VFS semantics and file system implementations introduces vulnerabilities.** The correctness of VFS is crucial to the underlying file system, however, it attracts much less attention than the integrity checking of the actual file system implementations. Operating systems often incorporate a Virtual File System (VFS), which offers uniform interfaces for user applications and redirects operations from applications to the underlying file system. For things to work correctly, the underlying file system must follow the specifications of the VFS. This is, however, often violated and brings many vulnerabilities. VFS vulnerabilities (such as CVE-2015-1420 and CVE-2015-2925) could affect a wide range of file system implementations. Many fault-tolerant file system designs, such as EnvyFS [13], also rely upon the correct operations of the VFS layer. However, existing work mostly focused on the bug detection for each individual file system (e.g., ext4, btrfs) [40, 60, 82], they rarely work on the correctness of the VFS layer [25]. Therefore, we believe applying bug detection techniques like model checking and fuzzing to VFS could also significantly improve file system security.

**F-8: Functionalities for supporting use-context can become the attack vectors of mobile and networked file systems.** Unlike generic file systems, mobile file systems and networked

file systems need special functionalities to support their use contexts. Implementation of these functionalities often involves vulnerabilities, which we detail below.

Mobile files systems typically run on portable devices like smart phones, which can easily get lost and encounter physical attacks. To protect user data under physical attacks, many mobile file systems introduce encryption mechanisms. For instance, APFS leverages AES-XTS or AES-CBC to encrypt user files on demand [1]. F2FS, instead of directly providing file encryption, utilizes a cryptographic stacked file system (eCryptFS) to encrypt data. However, vulnerabilities can arise in the implementation, often breaking the security of the encryption mechanisms. APFS leaks encryption keys via the Disk Utility hints (CVE-2017-7149) and APFS allows reading of decrypted data via Direct Memory Access (CVE-2017-13786). The eCryptFS on Samsung KNOX 1.0 uses a weak key generation algorithm, which makes brute-force attacks feasible (CVE-2016-1919).

Networked file systems need communication modules based on protocols like HTTP and RPC to support remote data transfer and access authentication. Many vulnerabilities appear in the communication modules, due to mishandling of malformed or crafted network packets, such as zero payloads, rafted packet header, and uninitialized packets. For example, GlusterFS allows RPC requests to create symbolic links (e.g., `gfs3_symlink_req`) pointing to file paths in the GlusterFS volume. This enables adversaries to send crafted RPC to create arbitrary symbolic links in the storage server and then execute arbitrary code (CVE-2018-10928).

**Summary:** Core components in file systems, due to their high complexities, contribute a majority of the vulnerabilities. The situation is worsening because new features are being integrated into the core components and are building up the complexities.

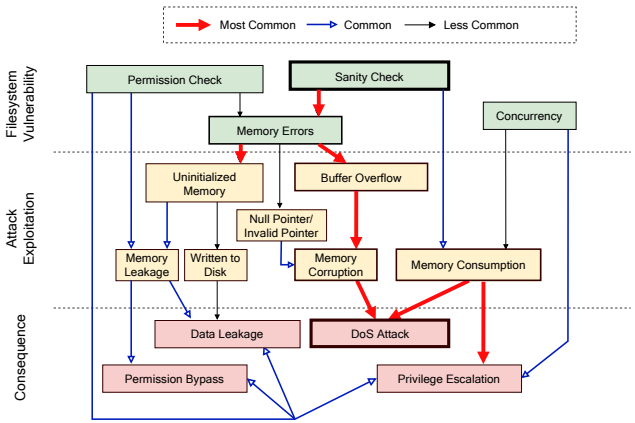
**Insight-2:** Different file system components introduce vulnerabilities in semantic-dependent ways. This suggests that different tools should be tailored, in a semantic-aware manner, to detect vulnerabilities in different components. For instance, to detect uninitialized memory in cache systems, a tool needs to identify not only explicit access (e.g., by regular file system operations) but also implicit access (e.g., by synchronization of cache management).

**Insight-3:** At present, security seems insufficiently considered when new file system features are added. It is desirable that the trade-off between functionality/efficiency and security is taken into account before new features are incorporated, particularly in security-sensitive sectors. References about which new features bring what security issues can be found from **F-3** to **F-8**. To fundamentally improve the situation, proactive measures, such as regression test [8, 57, 86] on the new features, must be taken.

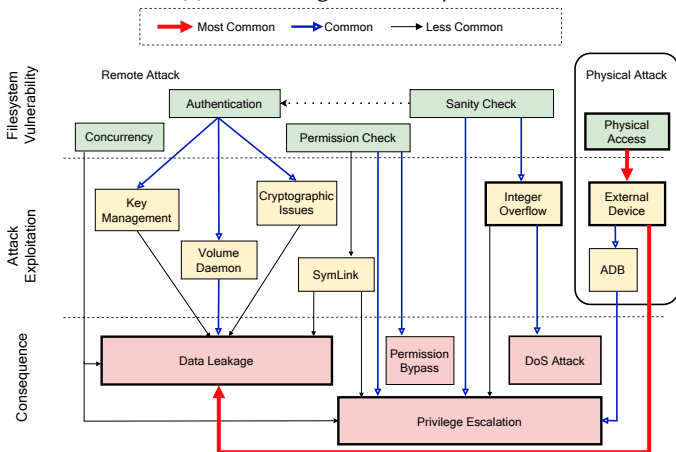
#### 4 THREATS OF FILE SYSTEM VULNERABILITIES

Our taxonomy brings insights towards mitigating vulnerabilities in file systems. However, file system vendors or distributors may still hesitate to take action because the threats of the vulnerabilities are unclear to them. In this section, we summarise the threats posed by file system vulnerabilities. We anticipate bringing evidence to motivate mitigation and bring insights into prioritizing the mitigation against more dangerous vulnerabilities.

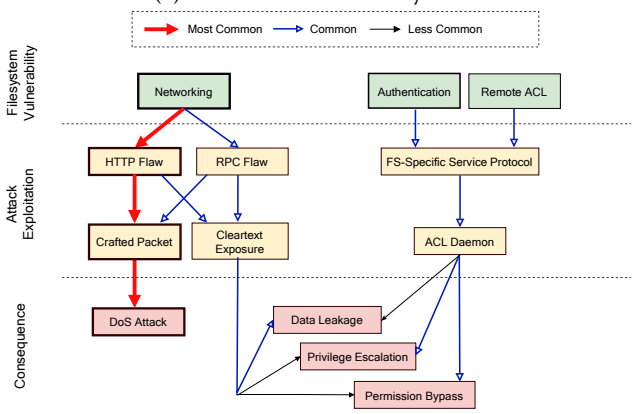
Overall, different vulnerabilities, when exploited, can lead to very diverse consequences. Figure 5 gives an overview of the trajectories from vulnerabilities in different file systems. The major attacks of generic file systems are caused by incomplete sanity checks which exploit memory errors like memory corruption and memory contention, which often result in Denial-of-Service (DoS). Mobile file systems have their unique attack paths, which mainly exploit the physical access vulnerabilities. They could bring severe consequences, such as data leakage. As networked file systems launches



(a) Threats to generic file systems



(b) Threats to mobile file systems



(c) Threats to networked file systems

Fig. 5. Trajectory from vulnerabilities to threats in different types of file systems.

Table 4. Distribution of file system vulnerabilities based on their threats. D: DoS, L: Data Leakage, E: Privilege Escalation, A: Access Permission Bypass, O: Other (e.g., write access to proc entries like CVE-2004-2613).

Type	Name	D	L	E	A	O
Generic	JFS [2]	50%	25%	0%	25%	0%
	Ext2 [17]	57%	14%	14%	14%	0%
	XFS [77]	70%	22%	4%	4%	0%
	VFS [4]	50%	4%	14%	25%	7%
	procfs [53]	44%	22%	22%	6%	6%
	ReiserFS [5]	60%	0%	40%	0%	0%
	Ext3 [76]	78%	11%	0%	11%	0%
	tmpfs [74]	33%	17%	50%	0%	0%
	Ext4 [50]	87%	4%	7%	2%	0%
	Btrfs [64]	54%	15%	8%	23%	0%
F2FS [42]	94%	0%	6%	0%	0%	
Mobile	HFS/HFS+ [36]	25%	13%	62%	0%	0%
	eCryptFS [34]	25%	25%	38%	0%	13%
	F2FS [42]	18%	55%	27%	0%	0%
	APFS [11]	0%	25%	75%	0%	0%
	Others	28%	17%	33%	11%	11%
Networked	NFS [59]	68%	1%	12%	18%	0%
	GlusterFS [27]	28%	17%	17%	38%	0%
	OpenAFS [3]	63%	30%	4%	4%	0%
	HDFS [35]	33%	17%	33%	17%	0%
	CephFS [79]	73%	20%	7%	0%	0%




Fig. 6. Vulnerability comparison of core OS components. In the figure, Severity, Impact, and Exploitability are extracted from the CVSS Score, the CIA Impacts, and the Access Complexity in CVE reports, respectively.

multiple file system instances, due to their incomplete packet filtering and validation, their network connection could be attacked with crafted packets via HTTP or RPC protocols. Table 4 presents the detailed distribution of vulnerabilities based on the threats they can pose. Following a similar organization of §3, we summarize our major findings below.

**F-9: File systems vulnerabilities impose a comparable level of security threat to vulnerabilities in other OS components.** Similar to our study on file system vulnerabilities, we gathered vulnerabilities in other operating system (OS) components that are committed to the CVE database in the past 20 years. As illustrated in Figure 6, file system vulnerabilities are as dangerous as vulnerabilities in other components, considering severity, impact, and exploitability as the metrics.

**F-10: DoS is the dominant threat brought by vulnerabilities to file systems.** About 61% of all the vulnerabilities can lead to DoS. They often first incur one of the following behaviors and eventually cause a DoS.

- *Memory corruption*: Memory errors and race conditions in file systems, as they typically do, can often cause memory corruption, followed by a kernel panic, such as the case CVE-2013-6382. Many semantic vulnerabilities can also lead to memory corruption and subsequent kernel panics. The previously discussed CVE-2015-8839 is such an example.
- *Memory consumption*: Another major cause of DoS in file systems is excessive memory consumption. Most of such cases are caused by memory leakage vulnerabilities. For instance, XFS before Linux 4.5.1 does not properly free the memory allocated for extended file attributes (CVE-2016-9685), which can be exploited to exhaust the memory and cause a DoS.
- *System hang*: CPU over-consumption, mainly brought by infinite loops and deadlocks, can turn into DoS. For instance, the integer overflow reported in CVE-2017-18257 can be exploited to trigger an infinite loop in the block allocation of F2FS, causing a system hang.
- *Network congestion*: Vulnerabilities derived from protocol flaws, which account for about 16% of all the vulnerabilities in networked file systems, frequently trigger DoS. As we pointed out before, such vulnerabilities typically happen because of mishandling of malformed network packets, which often easily disrupt normal communication and result in DoS. For example, the remote procedure call (RPC) module in OpenAFS (before 1.6.23 or 1.8.x before 1.8.2) does not restrict the size of an input. Adversaries can send, or claim to send, large inputs and consume server resources waiting for those inputs, blocking service to other valid connections (e.g., CVE-2018-16949).

**F-11: Data leakage is another big threat that vulnerabilities pose to file systems.** Three major types of file system vulnerabilities often lead to data leakage. The first type is uninitialized memory, where sensitive information, such as user data, file system metadata, and kernel memory unintentionally propagates to unauthorized destinations (e.g., CVE-2005-0400 and CVE-2004-0177). *In comparison to other software systems, uninitialized memory in file systems can more often lead to data leakage.* As we pointed out in §3, file systems incorporate many cache systems to improve their performance. Memory objects used by the cache systems shall be flushed to the disk, which creates extra paths for uninitialized memory to leak.

The second type is permission-bypassing vulnerabilities. By exploiting such vulnerabilities, the adversaries can gain access to files without the proper authorization and thus, indirectly leak private data. For instance, CVE-2012-4508 allows adversaries to obtain sensitive information from a deleted file by exploiting a race condition to gain permissions.

The last type is vulnerabilities in mobile file systems that enable physical attacks. The most direct consequence of such attacks is the breach of private data. For instance, physical adversaries, by exploiting CVE-2014-7951 in Android Debug Bridge of Android 4.0.4, can gain a direct connection to the device to read/write arbitrary files in the file system.

**F-12: Attackers can bypass access control via exploiting vulnerabilities in the enforcement of file permissions.** Flaws in the enforcement of permissions largely become vulnerabilities. In comparison to other types of vulnerabilities, permission-related vulnerabilities can often be reliably exploited and consistently lead to permission bypassing.

Beyond the authentication of user identities, networked file systems also run access control on file access requests. Two categories of vulnerabilities in this type of access control can also be exploited to gain file permissions on the removed servers: (1) validation operations before setting ACL are missing; (2) permission checks are missing. For example, the NFS client in Linux 2.6.29.3 or before misses a check on the permission bits for execution, allowing local users to bypass permissions and execute files on an NFSv4 fileserver (CVE-2009-1630). Similar issues are also reported in CVE-2005-3623 and CVE-2016-1237.

**F-13: Adversaries can achieve privilege escalation via exploiting file system vulnerabilities.** All the file systems we studied run at a high privilege. A significant subset of sanity-check vulnerabilities can enable adversaries to improperly “inherit” the high privilege. Some representative examples are: (i) CVE-2017-5551 preserves the `setgid` bit during a `setxattr` call in a `tmpfs` file system, allowing local users to gain group privileges with restrictions on execute permissions; (ii) CVE-2016-5393 wrongly gives HDFS service privileges to users who can only authenticate with the HDFS NameNode. This enables arbitrary commands execution with the same privileges as the HDFS service; (iii) CVE-2016-1572 does not validate mount destination file system types, which allows local users to gain privileges by mounting over a nonstandard file system.


Authentication protocol used in modern networked file systems can also fail due to vulnerabilities, and thus, give users undeserved privileges. Current networked file systems either develops its own authentication protocol [79] or reuses the third-party infrastructure [58, 75]. Both ways of authentication can be tricked and bypassed. We observe a large number of such cases in Ceph and HDFS, mostly caused by mishandling of user credential messages (e.g., CVE-2013-4134, CVE-2009-3516).

**F-14: Vulnerabilities in mobile file systems can sabotage the protection of data privacy.** To secure private data, mobile file systems have adopted protection mechanisms, mainly including encrypting user files [1] and isolating critical data inside trusted execution environments (TEEs). However, both mechanisms can be bypassed due to vulnerabilities. Since vulnerabilities related to file encryption have been discussed in §3 (see F-8), we now focus on TEEs below.


TEE has become a standard feature on mobile devices, which has been leveraged by mobile file systems to enable an isolated environment for protecting sensitive data. However, a drop-in deployment of TEE can be insufficient due to vulnerabilities in the interactions with TEE. For instance, a system crash in the non-secure domain of a TrustZone-enabled ARM platform will trigger an interrupt to switch the system context to the secure domain. Therein, CPU, memory, and register states can be accessed. However, an armored adversary can prevent the non-secure domain to issue the context switch, and initiate a Man-in-the-middle attack to obtain the execution information. For another example, eCryptFS uses an algorithm to generate a 32-byte AES key that combines the user password and 32-byte TIMA (TrustZone-based Integrity Measurement) key. The vulnerability (CVE-2016-1919) occurs because `Base64.getEncoder` expends the input with a ratio of 4:3. This means that only 24 bytes determine the final eCryptFS key and adversaries can easily crack the key with a brute-force search.

**F-15: Protocol flaws in networked file systems can enable code injection attacks.** Consider CephFS as an example. It uses the Ceph Object Gateway daemon (RGW) as an HTTP server to communicate with Ceph storage cluster and retrieve file data. The HTTP response splitting (HRS) flaw (e.g., CVE-2015-5245) in the RGW can enable the well-known Carriage Return Line Feed (CRLF) injection and cross-site scripting (XSS) attacks. Once the malicious code is injected, attackers will have the privilege to compromise the victim server, and then tamper the entire networked storage cluster.


**F-16: Vulnerabilities in a single instance of networked file system can impose threats to the entire system.** Networked file systems (e.g., HDFS [70] and GFS [26]) can often run a group of parallel instances on a cluster of servers. Each instance runs atop the local file system and delegates the file system to store data on the disk. The consequence of vulnerabilities in a single instance can propagate to all other instances in the same group. Consider GlusterFS as an example. GlusterFS supports *brick multiplexing* to reduce memory consumption, which allows multiple compatible bricks to share the same process and manage their own data volumes. However, a `NULL` pointer dereference (CVE-2018-10914) triggered by a client request in one instance can interrupt the brick



(a) The threats of generic file system vulnerabilities.



(b) The threats of mobile file system vulnerabilities.



(c) The threats of networked file system vulnerabilities.

Fig. 7. Evolutionary trend of the threats imposed by file system vulnerabilities. As for mobile file systems, no CVEs were reported before 2009.

process. This will propagate to other multiplexed brick processes, crushing the entire GlusterFS. Other similar cases include CVE-2012-4417, CVE-2019-15538, CVE-2020-24394.

**F-17: Vulnerabilities in different types of file systems present a different evolutionary trend in their threats.** As shown in Figure 7, vulnerabilities in *generic* file systems tend to result in more DoS and fewer attacks of other types. However, vulnerabilities in *mobile* file system and *networked* file systems present an opposite evolutionary trend. A presumable reason is vulnerabilities in generic file systems are more conventional and their consequences are limited by mitigation mechanisms such as memory sanitizer [67] and thread sanitizer [68]. In contrast, many vulnerabilities in mobile/networked file systems belong to newly emerging types (e.g., protocol flaws) and have been less mitigated. Considering all types of file systems together, the evolution in the threat of their vulnerabilities presents no clear trend.




Fig. 8. Distribution of file system patches, based on the lines of code (LoC) changed.

**Summary:** The threats of file system vulnerabilities, varying across different types, ranging from more functionality-critical ones like DoS to more security-critical ones like privilege escalation. History data shows that file system vulnerabilities, in general, are not becoming less harmful.

**Insight-4:** Our study unveils that different file system operations, when becoming vulnerable, can turn into different threats. This brings insights towards the mitigating process (e.g., patching): vulnerabilities involved in operations that can result in higher threats shall be prioritized. For instance, uninitialized memory errors in cache systems, which often incur data leakage, deserve an earlier fix than excessive memory consumption that typically incurs DoS. References about which file system operations can lead to what threats can be found from F-10 to F-16.

**Insight-5:** Our study also unveils that vulnerabilities in mobile/networked file systems are presenting increasing threats (e.g., more data leakage than DoS). Considering the increasing popularity of mobile/networked file systems in the coming era of IoT, our study brings evidence that more efforts should be prioritized to mitigate vulnerabilities in mobile/networked file systems.

## 5 PATCHING OF FILE SYSTEM VULNERABILITIES

Our taxonomy and threat analysis offer insights and motivations to mitigate file system vulnerabilities. We take a step further to analyze their patches and gain deeper insights. This will be beneficial to file system developers, since it can help people learn about the common patterns of patching file systems and guide them to avoid the deficiencies in today's practice. In the patch analysis process, we aim to answer three questions: (1) *how complex are the patches?* (2) *how are the patches applied, systematically or in an ad-hoc manner?* (3) *are patches across different file systems correlated, and how that may affect the patching process?*

**F-18: Most of the patches (56%) produce small modifications to the file systems.** Figure 8 shows the distribution of patches based on their complexities (i.e., how many lines of code are modified). We find that a large number of patches are simply adding a check to the memory boundary (e.g., CVE-2013-6382) or fixing a syntax error (e.g., CVE-2013-1848). These patches constitute the low-complexity group. Most of the strategies are conventional, such as adding boundary checks to fix buffer overflows (e.g., CVE-2008-3531). For instance, the patches of many race conditions simply make the responsible operations atomic. The drawback of such ad-hoc strategies is they are too passive. They can neither help prevent similar vulnerabilities in the codebase nor bring global benefits like new designs. We believe a more systematic strategy should be taken (see our discussion in F-20).

**F-19: A non-negligible portion (20%) of the patches involve complex modifications, such as fixes to multiple file systems and comprehensive validations of file system data structures. Some vulnerabilities that are shared across file systems can cause duplicated patches.**



---

```
1 + int setattr_killpriv(struct dentry *dentry, struct iattr *iattr)
```

---

(a) A new interface added to VFS that removes extended privilege attributes.

---

```
1 int ext3_setattr(struct dentry *dentry,
2     struct iattr *iattr)
3     if (error)
4         return error;
5 + error = setattr_killpriv(dentry, iattr);
6 + if (error)
7 + return error;
```

---

(b) Patch to ext3 filesystem.

---

```
1 int btrfs_setattr(struct dentry *dentry,
2     struct iattr *iattr)
3     if (err)
4         return err;
5 + err = setattr_killpriv(dentry, iattr);
6 + if (err)
7 + return err;
```

---

(c) Patch to btrfs filesystem.

Fig. 9. Similar patches to different Linux file systems to reflect the same VFS inference update.

In most of the patches we studied, the patches aim to solve a unique flaw in a certain file system. In the study, we discovered that for a shared vulnerability in multiple file systems, developers release patches that have a similar strategy or even the same code logic. These patches build up the high-complexity group. Strategies of a certain patch to a specific file system can be borrowed by patches that are used by other types of file systems. For example, as shown in Figure 9a, the VFS subsystem added a new function `setattr_killpriv` to perform additional permission and sanity checks for the extended attributes (xattrs). Since most Linux file systems supports xattrs, they have to properly invoke this function in their own `setattr()` to ensure security. As a result, these file systems generate similar patches (Figure 9b and c, CVE-2015-1350). Such pattern also occurs in the patching of `xattr` block caching in the ext2 and ext4 file systems. Both of the file systems have to convert the original meta block cache from `mbcache` to `mbcache2`. Hence, they should add the same sanity checking to the entries and generate similar patches (e.g., CVE-2015-8952).

A cross-filesystem patching approach could offer two major benefits. First, having a unified solution simplifies the patch development process as well as the future maintenance. Second, file systems which share the patch can take advantage of further improvements in each individual file system, which enables us to improve the overall robustness and reliability.

**F-20: Some patches demonstrate systematic strategies, which can benefit vulnerability prevention.** Despite ad-hoc strategies still being dominant, we do observe systematic strategies in two cases. First, there are patches that perform systematic verification on file system data structures instead of simply applying an ad-hoc fix. For instance, CVE-2018-14613 reports a simple invalid pointer dereference issue when mounting and operating a crafted `btrfs` image caused by invalid block group items. This specific vulnerability could have been solved by a low complexity patch. However, the patch to it demonstrates good practice by adding a new sanity check function that performs a complete check of the item size, offset, object id, type, and used space. This not only helps fix the reported vulnerabilities, but also removes the undiscovered hidden flaws to prevent future vulnerabilities. Indeed, none of the follow-up CVEs has related issues when mounting and operating on `btrfs` images. Second, many patches of race conditions switch from lock-based implementations to lock-free implementations, such as deploying semaphore-based solutions (e.g., CVE-2014-9710) or avoiding race conditions in the first place (e.g., CVE-2014-8086). These patches bring design-wise insights, which are applicable to a more fundamental resolution of race conditions.

Table 5. Suggested strategies to patch common types of file system vulnerabilities. The third column shows the percentage of CVEs that can be fixed with the suggested strategies.

Type	Suggested Defense Strategies <sup>1</sup>	% of Applicable CVEs
Sanity checks	FS image validation before mounting *	48.2%
	Validating fs parameters upon file access	12.5%
	Memory pointer validation and bounds checking	5.4%
	Handle exceptions of division by zero	3.9%
	Avoid integer overflows	3.6%
Conc. issues	Enforce atomic ToCToU race operations	42.9%
	Support of lock-free optimizations *	28.6%
Memory errors	Initialize unused, allocated memory	38.5%
	Initialize data before copying them from user space	24.2%
	Reinitialize fs metadata upon file updates	14.3%
File pems.	Check enforcement of user privileges	42.9%
	Check file permissions for direct storage access	28.6%
	Validate fs metadata information before mounting fs	4.8%

<sup>1</sup> \* indicates the mitigation is systematic.

**F-21: Patches that modify the VFS need coordinated updates to the underlying file systems.** A patch to the VFS subsystem often requires updating the underlying file systems. Otherwise, the patch may not come into effect. In CVE-2015-1350, the VFS subsystem provides an incomplete set of requirements for `setattr` operations, which under-specifies the removal of extended privilege attributes. The corresponding patch silently introduces a new VFS API to fix the flaw. Concurrently, the patch has to modify all 21 affected underlying file systems to use the new API. While, in this example, the complexity incurred by the need for coordinated patching is well handled, it is not always so. For example, the patch to CVE-2016-7097 modifies VFS but fails to update the complete set of 15 affected underlying file systems, leaving the vulnerability still open (reported as CVE-2017-5551). Until one year later, another patch to CVE-2017-5551 eventually fixed the problem.

**Summary:** Patches to file system vulnerabilities carry different complexities. The patches also adopt various strategies. Most of them are ad-hoc, bringing no help with addressing similar but yet unknown vulnerabilities. However, patches following systematic strategies are arising. These patches present preventive measures or new designs, offering insights towards fundamentally mitigating vulnerabilities of the same root causes. Finally, file systems are often correlated. However, today's practice often overlooks the correlated file systems when patching the target ones, thus leading to insecure patches.

**Insight-6:** Our study unveils that the current practice of patching file systems is largely passive. Continuing this practice will likely prolong the arm race between vulnerability creation and patching. To fundamentally escalate the security of file systems, the community should encourage and promote the adoption of preventive and secure-by-design patching. Exemplary strategies observed by our study have been presented in **F-20**.

**Insight-7:** Better attention needs to be drawn on the correlation of file systems when developing patches. To ensure the correctness of patches, any modification to a file system (in particular VFS) must be validated to understand its impacts on other file systems. Whenever such impacts arise, the affected file systems must be updated accordingly.

## 6 SUGGESTIONS

With this study, we discover a group of vulnerabilities that **our community is not handling well at this point**. Most of them are semantic related or use-context related. We summarize the

Table 6. Suggested defense strategies for mobile *fs* vulnerabilities. The second column shows the percentage of CVEs that could be fixed with the suggested strategies.

Mobile Defense Strategies	% of Applicable CVEs
Enhance authentication mechanisms in ADB channels	25.0%
Enhance USB debugging mode	18.8%
Enforce secure interaction between TrustZone and FS	12.5%
(Total)	56.3%

Table 7. Suggested defense strategies for networked *fs* vulnerabilities. The second column shows the percentage of CVEs that could be fixed with the suggested strategies.

Unique Defense Strategies	% of Applicable CVEs.
Enforce strict authentication/permission control	26.4%
Verify (HTTP/RPC) network packets	23.1%
Enforce VFS specifications	9.3%
Fault isolation between <i>fs</i> instances	7.1%
(Total)	65.9%

common suggested strategies used by the patches on generic file system vulnerabilities in Table 5. The common defense strategies discussed of patching generic file system can be also applied to fix 43.7% of the mobile file system and 34.4% of the networked file system vulnerabilities. Meanwhile, we identify their unique defense strategies vulnerabilities, as shown in Table 6 and 7.

**S-1: Verifying file system implementations with semantics specified in VFS.** As discussed, VFS provides a uniform abstraction level for upper-level programs to interact with real file system implementations. It offers semantics (e.g., POSIX) for accessing the underlying file systems. However, the real filesystem implementation may not exactly satisfy the semantics, due to the unclear specifications between the VFS and low-level filesystem implementations. Given that VFS specifications have been defined, they can be refined to verify the filesystem implementations. Prior verification studies of file systems [55, 71] mostly worked on the crash safety. Our findings reveal that more efforts are required to achieve the end-to-end verification of each semantic mapping from the high-level VFS specification to the low-level filesystem implementation.

**S-2: Enhancing secure channels between mobiles and external physical devices.** To defend against physical attacks, we need to continue developing secure channels between mobile devices and external devices such as Thunderbolt adapters and USB ports. According to our study, nearly 23% of vulnerabilities in Android can be exploited by evil-maid attacks. These attacks can be prevented by explicitly authorizing the mounting of external devices.

**S-3: Enforcing secure communication between TEE and file systems.** We still cannot solely rely on TEE or encrypted file systems or even both to ensure the security of mobile file systems. Particularly, we have to ensure the secure interactions between TEE and file systems, such as preventing the direct access to the memory used by Volume Daemon or storage management process from TrustZone. We have to carefully refine the interfaces between them, which can be improved with the development at both TEE and file system side [19, 61, 81].

**S-4: Verifying network packet.** According to our study, verifying the network packets, especially those via the HTTP or RPC protocols, can significantly reduce the risk of network attacks in networked file systems. The semantic specifications from upper-level service protocols defined by network file systems will facilitate such network packet verification. Similar network verification techniques have been developed in software-defined networking domain [62, 63] and packet level authentication [45], however, few of them focused on the semantic-aware verification of storage traffic from networked file systems.

**S-5: Fault isolation for networked filesystem instances.** As discussed in §4, a filesystem instance (e.g., a data volume manager) has correlations with many system components such as local file systems and OS kernel. These correlations significantly complicate the management of filesystem instances and increase the attack surface. Providing a fault isolation mechanism for filesystem instances can significantly enhance the security of the entire networked system. Systems techniques, such as containers [44, 69], sandbox [16], and TEE [12], have been proposed to enforce the isolation between programs. They can be leveraged to enable isolation between filesystem instances. However, a holistic approach that requires fine-grained function partition and placement for networked file systems is still highly desirable [87].

## 7 RELATED WORK

**Study of Bugs and Vulnerabilities.** Past research has conducted many studies on defects in operating systems. The studies vary in targets. Many of them [47, 52, 78, 82] focus on functionality bugs in file systems while some others [28, 38] concentrate on virtual memory management problems. The studies also vary in scope that ranges from concurrent issues [24, 32, 43, 49] to specific families of vulnerabilities [21, 46, 66]. Our work uses a methodology similar to many of the studies. However, we focus on the security vulnerabilities in file systems. Unlike the previous file system studies that focused on bug causes and consequences [15, 47], our study centers around the understanding of the attack surface of file systems and the vulnerability exploitation procedure, while covering a wider range of CVEs across different types of file systems. To the best of our knowledge, our study is the first in-depth work of its kind.

**Bug and Vulnerability Detection.** Research in this line can be classified into two categories: bug finding and formal verification. Bug finding tools, such as FiSC [85], EXPLODE [84], and JUXTA [52], can discover file system bugs based on semantic-aware patterns. Our study can benefit these tools by providing insights into expanding their patterns. Formal verification [10, 14, 20, 22, 30, 55, 71] is also promising. However, at this stage, it is still challenging to verify the security of an entire file system due to the high complexity [30, 71, 80]. Our study can complement formal verification. It pinpoints the file system components that are more vulnerable, thus enabling formal verification to narrow down the scope and scale up.

**Secure File Systems.** Besides bug study and detection, researchers have also been endeavoring to build secure and reliable file systems [48, 51]. Lu et al. [48] proposed the physical disentanglement to minimize storage faults propagation. Min et al. [51] leveraged transaction flash storage to build crash-consistent file systems. Our study can bring insights to follow-up works in this line of research. For instance, we find that many vulnerabilities in networked file systems are caused by the lack of fault isolation across different instances. This can motivate efforts to develop fault isolation mechanisms for file systems.

## 8 CONCLUSION

This paper presents an empirical study on the security of modern file systems, following the angle of inspecting vulnerabilities disclosed from mainstream file systems. Throughout the study, we build the first systematic understanding of the attack surfaces faced by file systems, the threats tied to the attack surfaces, and the limitations in today’s practice of addressing the attack surfaces. We envision that our study will raise awareness of file system security and, more importantly, offer insights towards improving file system security.

## REFERENCES

- [1] Apple File System secure encryption mechanism. <https://9to5mac.com/2016/06/13/apple-file-system-apfs/>.
- [2] Journaled File System Technology for Linux. <http://jfs.sourceforge.net>.
- [3] OpenAFS. <https://www.openafs.org>.
- [4] Overview of the Linux Virtual File System. <https://www.kernel.org/doc/html/latest/filesystems/vfs.html>.
- [5] ReiserFS. <https://en.wikipedia.org/wiki/ReiserFS>.
- [6] Metadata replication for ext4. <https://lwn.net/Articles/465041/>, 2011.
- [7] Request cve ids. [https://cve.mitre.org/cve/request\\_id.html](https://cve.mitre.org/cve/request_id.html), 2021.
- [8] Hiralal Agrawal, Joseph Robert Horgan, Edward W Krauser, and Saul A London. Incremental regression testing. In *1993 Conference on Software Maintenance*, pages 348–357. IEEE, 1993.
- [9] Adil Ahmad, Kyungtae Kim, Muhammad Ihsanulhaq Sarfaraz, and Byoungyoung Lee. Obliviate: A data oblivious filesystem for intel sgx. In *Network and Distributed System Security Symposium, (NDSS '18)*, San Diego, CA, USA, 2018.
- [10] Sidney Amani, Alex Hixon, Zilin Chen, Christine Rizkallah, Peter Chubb, Liam O'Connor, Joel Beeren, Yutaka Nagashima, Japheth Lim, Thomas Sewell, Joseph Tuong, Gabriele Keller, Toby C. Murray, Gerwin Klein, and Gernot Heiser. Cogent: Verifying high-assurance file system implementations. In *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'16)*, Atlanta, GA, 2016.
- [11] APFS: Apple File System. [https://en.wikipedia.org/wiki/Apple\\_File\\_System](https://en.wikipedia.org/wiki/Apple_File_System).
- [12] Pierre-Louis Aublin, Florian Kelbert, Dan O'Keeffe, Divya Muthukumar, Christian Priebe, Joshua Lind, Robert Krahn, Christof Fetzer, David Eyers, and Peter Pietzuch. Libseal: Revealing service integrity violations using trusted execution. In *Proceedings of the Thirteenth EuroSys Conference (EuroSys'18)*, 2018.
- [13] Lakshmi N Bairavasundaram, Swaminathan Sundararaman, Andrea C Arpaci-Dusseau, and Remzi H Arpaci-Dusseau. Tolerating file-system mistakes with envyfs. In *USENIX Annual Technical Conference*, volume 9, 2009.
- [14] James Bornholt, Antoine Kaufmann, Jialin Li, Arvind Krishnamurthy, Emina Torlak, and Xi Wang. Specifying and checking file system crash-consistency models. In *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '16*, Atlanta, GA, USA, April 2016.
- [15] Miao Cai, Hao Huang, and Jian Huang. Understanding security vulnerabilities in file systems. *APSys '19*, page 8–15, New York, NY, USA, 2019. Association for Computing Machinery.
- [16] Justin Capps, Armon Dadgar, Jeff Rasley, Justin Samuel, Ivan Beschastnikh, Cosmin Barsan, Arvind Krishnamurthy, and Thomas Anderson. Retaining sandbox containment despite bugs in privileged memory-safe code. In *Proceedings of the 17th ACM Conference on Computer and Communications Security (CCS'10)*, 2010.
- [17] Remy Card. Design and implementation of the second extended filesystem. In *Proceedings of the First Dutch International Symposium on Linux, 1995*, 1995.
- [18] Ceph. [https://en.wikipedia.org/wiki/Ceph\\_\(software\)](https://en.wikipedia.org/wiki/Ceph_(software)).
- [19] David Cerdeira, Nuno Santos, Pedro Fonseca, and Sandro Pinto. Sok: Understanding the prevailing security vulnerabilities in trustzone-assisted tee systems. In *Proceedings of IEEE Symposium on Security and Privacy (Oakland'20)*, 2020.
- [20] Haogang Chen, Tej Chajed, Alex Konradi, Stephanie Wang, Atalay Ileri, Adam Chlipala, M. Frans Kaashoek, and Nikolai Zeldovich. Verifying a high-performance crash-safe file system using a tree specification. In *Proceedings of the 26th Symposium on Operating Systems Principles (SOSP '17)*, Shanghai, China, October 2017.
- [21] Haogang Chen, Yandong Mao, Xi Wang, Dong Zhou, Nikolai Zeldovich, and M. Frans Kaashoek. Linux kernel vulnerabilities: state-of-the-art defenses and open problems. In *APSys '11 Asia Pacific Workshop on Systems*, Shanghai, China, July 2011.
- [22] Haogang Chen, Daniel Ziegler, Tej Chajed, Adam Chlipala, M. Frans Kaashoek, and Nikolai Zeldovich. Using crash hoare logic for certifying the FSCQ file system. In *Proceedings of the 25th Symposium on Operating Systems Principles, SOSP 2015*, Monterey, CA, USA, October 2015.
- [23] Ext4 Disk Layout. [https://ext4.wiki.kernel.org/index.php/Ext4\\_Disk\\_Layout#Flexible\\_Block\\_Groups](https://ext4.wiki.kernel.org/index.php/Ext4_Disk_Layout#Flexible_Block_Groups), 2020.
- [24] Pedro Fonseca, Cheng Li, Vishal Singhal, and Rodrigo Rodrigues. A study of the internal and external effects of concurrency bugs. In *Proceedings of the 2010 IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2010*, Chicago, IL, USA, June 2010.
- [25] Andy Galloway, Gerald Lüttgen, Jan Tobias Mühlberg, and Radu I Siminiceanu. Model-checking the linux virtual file system. In *International Workshop on Verification, Model Checking, and Abstract Interpretation*, pages 74–88. Springer, 2009.
- [26] Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung. The google file system. 2003.
- [27] GlusterFS. Glusterfs architecture. <https://docs.gluster.org/en/v3/Quick-Start-Guide/Architecture/>.

- [28] Sudhakar Govindavajhala and Andrew W Appel. Using memory errors to attack a virtual machine. In *2003 Symposium on Security and Privacy (SP)*, Berkeley, CA, USA, 2003.
- [29] Andreas Gruenbacher. POSIX access control lists on linux. In *Proceedings of the 2003 USENIX Annual Technical Conference (USENIX ATC'03)*, San Antonio, TX, June 2003.
- [30] Ronghui Gu, Zhong Shao, Hao Chen, Xiongnan (Newman) Wu, Jieung Kim, Vilhelm Sjöberg, and David Costanzo. Certikos: An extensible architecture for building certified concurrent OS kernels. In *12th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2016*, Savannah, GA, uSA, November 2016.
- [31] Researcher Reservation Guidelines. [https://cve.mitre.org/cve/researcher\\_reservation\\_guidelines](https://cve.mitre.org/cve/researcher_reservation_guidelines), 2019.
- [32] Haryadi S. Gunawi, Mingzhe Hao, Tanakorn Leesatapornwongsa, Tiratat Patana-anake, Thanh Do, Jeffry Adityatama, Kurnia J. Eliazar, Agung Laksono, Jeffrey F. Lukman, Vincentius Martin, and Anang D. Satria. What bugs live in the cloud? A study of 3000+ issues in cloud systems. In *Proceedings of the ACM Symposium on Cloud Computing, SoCC 2014*, Seattle, WA, USA, November 2014.
- [33] Haryadi S Gunawi, Vijayan Prabhakaran, Swetha Krishnan, Andrea C Arpaci-Dusseau, and Remzi H Arpaci-Dusseau. Improving file system reliability with i/o shepherding. In *Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles*, pages 293–306, 2007.
- [34] Michael Austin Halcrow. ecryptfs: An enterprise-class encrypted filesystem for linux. In *Proceedings of the 2005 Linux Symposium*, volume 1, pages 201–218, 2005.
- [35] Hadoop Distributed File System. [https://hadoop.apache.org/docs/r1.2.1/hdfs\\_design.html](https://hadoop.apache.org/docs/r1.2.1/hdfs_design.html), 2019.
- [36] HFS: Hierarchical File System. [https://en.wikipedia.org/wiki/Hierarchical\\_File\\_System](https://en.wikipedia.org/wiki/Hierarchical_File_System).
- [37] Dave Hitz, James Lau, and Michael A. Malcolm. File system design for an NFS file server appliance. In *USENIX Winter 1994 Technical Conference*, San Francisco, CA, USA, January 1994.
- [38] Jian Huang, Moinuddin K. Qureshi, and Karsten Schwan. An Evolutionary Study of Linux Memory Management for Fun and Profit. In *Proceedings of the 2016 USENIX Annual Technical Conference (USENIX ATC'16)*, Denver, Co, June 2016.
- [39] Sudarsun Kannan, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, Yuangang Wang, Jun Xu, and Gopinath Palani. Designing a true direct-access file system with devfs. In *16th USENIX Conference on File and Storage Technologies (FAST'18)*, Oakland, CA, 2018.
- [40] Seulbae Kim, Meng Xu, Sanidhya Kashyap, Jungyeon Yoon, Wen Xu, and Taesoo Kim. Finding semantic bugs in file systems with an extensible fuzzing framework. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles (SOSP'19)*, Ontario, Canada, 2017.
- [41] Jinhung Koo, Junsu Im, Jooyoung Song, Juhung Park, Eunji Lee, Bryan S Kim, and Sungjin Lee. Modernizing file system through in-storage indexing. In *15th USENIX Symposium on Operating Systems Design and Implementation (OSDI'21)*, pages 75–92, 2021.
- [42] Changman Lee, Dongho Sim, Joo Young Hwang, and Sangyeun Cho. F2FS: A new file system for flash storage. In *Proceedings of the 13th USENIX Conference on File and Storage Technologies, FAST 2015*, Santa Clara, CA, USA, February 2015.
- [43] Tanakorn Leesatapornwongsa, Jeffrey F. Lukman, Shan Lu, and Haryadi S. Gunawi. Taxdc: A taxonomy of non-deterministic concurrency bugs in datacenter distributed systems. In *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'16)*, Atlanta, GA, April 2016.
- [44] Xin Lin, Lingguang Lei, Yewu Wang, Jiwu Jing, Kun Sun, and Quan Zhou. A measurement study on linux container security: Attacks and countermeasures. In *Proceedings of the 34th Annual Computer Security Applications Conference (ACSAC'18)*, 2018.
- [45] Jed Liu, William Hallahan, Cole Schlesinger, Milad Sharif, Jeongkeun Lee, Robert Soulé, Han Wang, Călin Cașcaval, Nick McKeown, and Nate Foster. P4v: Practical verification for programmable data planes. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication (SIGCOMM'18)*, 2018.
- [46] Kangjie Lu, Marie-Therese Walter, David Pfaff, Stefan Nüumberger, Wenke Lee, and Michael Backes. Unleashing use-before-initialization vulnerabilities in the linux kernel using targeted stack spraying. In *Network and Distributed System Security Symposium, (NDSS '17)*, San Diego, CA, USA, 2017.
- [47] Lanyue Lu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, and Shan Lu. A study of linux file system evolution. In *Proceedings of the 11th USENIX conference on File and Storage Technologies (FAST'13)*, San Jose, CA, 2013.
- [48] Lanyue Lu, Yupu Zhang, Thanh Do, Samer Al-Kiswany, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. Physical disentanglement in a container-based file system. In *11th USENIX Symposium on Operating Systems Design and Implementation, OSDI '14*, Broomfield, CO, USA, October 2014.
- [49] Shan Lu, Soyeon Park, Eunsoo Seo, and Yuanyuan Zhou. Learning from mistakes: a comprehensive study on real world concurrency bug characteristics. In *Proceedings of the 13th International Conference on Architectural Support for*

- Programming Languages and Operating Systems, ASPLOS 2008*, Seattle, WA, USA, March 2008.
- [50] Avantika Mathur, Mingming Cao, Suparna Bhattacharya, Andreas Dilger, Alex Tomas, and Laurent Vivier. The new ext4 filesystem: current status and future plans. In *Proceedings of the Linux symposium, 2007*.
  - [51] Changwoo Min, Woon-Hak Kang, Taesoo Kim, Sang-Won Lee, and Young Ik Eom. Lightweight application-level crash consistency on transactional flash storage. In *2015 USENIX Annual Technical Conference, USENIX ATC '15*, Santa Clara, CA, USA, July 2015.
  - [52] Changwoo Min, Sanidhya Kashyap, Byoungyoung Lee, Chengyu Song, and Taesoo Kim. Cross-checking semantic correctness: the case of finding file system bugs. In *Proceedings of the 25th Symposium on Operating Systems Principles, SOSP 2015*, Monterey, CA, USA, October 2015.
  - [53] Erik Mouw. Linux kernel procs guide. *Delft University of Technology*, 2001.
  - [54] Dongliang Mu, Alejandro Cuevas, Limin Yang, Hang Hu, Xinyu Xing, Bing Mao, and Gang Wang. Understanding the reproducibility of crowd-reported security vulnerabilities. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*, pages 919–936, 2018.
  - [55] Luke Nelson, Helgi Sigurbjarnarson, Kaiyuan Zhang, Dylan Johnson, James Bornholt, Emina Torlak, and Xi Wang. Hyperkernel: Push-button verification of an os kernel. In *Proceedings of the 26th Symposium on Operating Systems Principles (SOSP'17)*, 2017.
  - [56] Network File System.  
[https://en.wikipedia.org/wiki/Network\\_File\\_System](https://en.wikipedia.org/wiki/Network_File_System), 2020.
  - [57] Akira K Onoma, Wei-Tek Tsai, Mustafa Poonawala, and Hiroshi Suganuma. Regression testing in an industrial environment. *Communications of the ACM*, 41(5):81–86, 1998.
  - [58] Brian Pawlowski, Chet Juszczak, Peter Staubach, Carl Smith, Diane Lebel, and Dave Hitz. Nfs version 3: Design and implementation. In *USENIX Summer Technical Conference*, Boston, MA, USA, 1994.
  - [59] Brian Pawlowski, David Noveck, David Robinson, and Robert Thurlow. The nfs version 4 protocol. In *In Proceedings of the 2nd International System Administration and Networking Conference (SANE 2000)*, Maastricht, Netherlands, 2000.
  - [60] Thanumalayan Sankaranarayanan Pillai, Vijay Chidambaram, Ramnatthan Alagappan, Samer Al-Kiswany, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. All file systems are not created equal: On the complexity of crafting crash-consistent applications. In *11th USENIX Symposium on Operating Systems Design and Implementation, OSDI '14*, Broomfield, CO, USA, October 2014.
  - [61] Sandro Pinto and Nuno Santos. Demystifying arm trustzone: A comprehensive survey. *ACM Comput. Surv.*, 51(6), January 2019.
  - [62] Gordon D. Plotkin, Nikolaj Børner, Nuno P. Lopes, Andrey Rybalchenko, and George Varghese. Scaling network verification using symmetry and surgery. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'16)*, 2016.
  - [63] Santhosh Prabhu, Kuan Yen Chou, Ali Kheradmand, Brighten Godfrey, and Matthew Caesar. Plankton: Scalable network configuration verification through model checking. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI'20)*, Santa Clara, CA, February 2020.
  - [64] Ohad Rodeh, Josef Bacik, and Chris Mason. Btrfs: The linux b-tree filesystem. *ACM Transactions on Storage (TOS)*, 9(3):1–32, 2013.
  - [65] Mendel Rosenblum and John K. Ousterhout. The design and implementation of a log-structured file system. In *Proceedings of the Thirteenth ACM Symposium on Operating System Principles, SOSP 1991*, Pacific Grove, CA, USA, October 1991.
  - [66] Sergej Schumilo, Cornelius Aschermann, Robert Gawlik, Sebastian Schinzel, and Thorsten Holz. k afl: Hardware-assisted feedback fuzzing for {OS} kernels. In *26th {USENIX} Security Symposium ({USENIX} Security 17)*, San Jose, CA, USA, 2017.
  - [67] Konstantin Serebryany, Derek Bruening, Alexander Potapenko, and Dmitriy Vyukov. Addresssanitizer: A fast address sanity checker. In *2012 {USENIX} Annual Technical Conference ({USENIX} {ATC} 12)*, pages 309–318, 2012.
  - [68] Konstantin Serebryany and Timur Iskhodzhanov. Threadsanitizer: Data race detection in practice. In *Proceedings of the workshop on binary instrumentation and applications*, pages 62–71, 2009.
  - [69] Zhiming Shen, Zhen Sun, Gur-Eyal Sela, Eugene Bagdasaryan, Christina Delimitrou, Robbert Van Renesse, and Hakim Weatherspoon. X-containers: Breaking down barriers to improve performance and isolation of cloud-native containers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'19)*, 2019.
  - [70] Konstantin Shvachko, Hairong Kuang, Sanjay Radia, Robert Chansler, et al. The hadoop distributed file system. In *IEEE 26th Symposium on Mass Storage Systems and Technologies, (MSST '10)*, Incline Village, NV, USA, 2010.
  - [71] Helgi Sigurbjarnarson, James Bornholt, Emina Torlak, and Xi Wang. Push-button verification of file systems via crash refinement. In *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI'16)*, Savannah, GA, USA, November 2016.

- [72] Satyam Singh. Exploiting nfs share. <https://resources.infosecinstitute.com>, 2021.
- [73] Stephen Smalley, Chris Vance, and Wayne Salamon. Implementing selinux as a linux security module. *NAI Labs Report*, 1(43):139, 2001.
- [74] Peter Snyder. tmpfs: A virtual memory file system. In *Proceedings of the autumn 1990 EUUG Conference*, pages 241–248, 1990.
- [75] Jennifer G Steiner, B Clifford Neuman, and Jeffrey I Schiller. Kerberos: An authentication service for open network systems. In *Usenix Winter Technical Conference*, Dallas, TX, USA, 1988.
- [76] Stephen Tweedie. Journaling the linux ext2fs filesystem. In *The Fourth Annual Linux Expo, 1998*, Durham, NC, USA, May 1998.
- [77] Randolph Y Wang and Thomas E Anderson. xfs: A wide area mass storage file system. In *Proceedings of IEEE 4th Workshop on Workstation Operating Systems. WWOS-III*, pages 71–78. IEEE, 1993.
- [78] Jinpeng Wei and Calton Pu. Tocttough vulnerabilities in unix-style file systems: An anatomical study. In *4th USENIX Conference on File and Storage Technologies, (FAST '05)*, San Francisco, CA, USA, December 2005.
- [79] Sage A Weil, Scott A Brandt, Ethan L Miller, Darrell DE Long, and Carlos Maltzahn. Ceph: A scalable, high-performance distributed file system. In *Proceedings of the 7th symposium on Operating systems design and implementation (OSDI'06)*, Seattle, WA, USA, 2006.
- [80] Fengwei Xu, Ming Fu, Xinyu Feng, Xiaoran Zhang, Hui Zhang, and Zhaohui Li. A practical verification framework for preemptive OS kernels. In *Computer Aided Verification - 28th International Conference, CAV 2016*, Toronto, ON, Canada, July 2016.
- [81] Meng Xu, Chengyu Song, Yang Ji, Ming-Wei Shih, Kangjie Lu, Cong Zheng, Ruian Duan, Yeongjin Jang, Byoungyoung Lee, Chenxiong Qian, Sangho Lee, and Taesoo Kim. Toward engineering a secure android ecosystem: A survey of existing techniques. *ACM Comput. Surv.*, 49(2), August 2016.
- [82] Wen Xu, Hyungon Moon, Sanidhya Kashyap, Po-Ning Tseng, and Taesoo Kim. Fuzzing file systems via two-dimensional input space exploration. In *Proceedings of the 2019 IEEE Symposium on Security and Privacy (Oakland'19)*, San Francisco, CA, USA, 2019.
- [83] Fabian Yamaguchi, Nico Golde, Daniel Arp, and Konrad Rieck. Modeling and discovering vulnerabilities with code property graphs. In *2014 IEEE Symposium on Security and Privacy*, pages 590–604. IEEE, 2014.
- [84] Junfeng Yang, Can Sar, and Dawson R. Engler. EXPLODE: A lightweight, general system for finding serious storage system errors. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation (OSDI'06)*, Seattle, WA, USA, November 2006.
- [85] Junfeng Yang, Paul Twohey, Dawson R. Engler, and Madanlal Musuvathi. Using model checking to find serious file system errors. In *6th Symposium on Operating System Design and Implementation (OSDI 2004)*, San Francisco, CA, USA, December 2004.
- [86] Shin Yoo and Mark Harman. Regression testing minimization, selection and prioritization: a survey. *Software testing, verification and reliability*, 22(2):67–120, 2012.
- [87] Ennan Zhai, Ruichuan Chen, David Isaac Wolinsky, and Bryan Ford. Heading off correlated failures through independence-as-a-service. In *Proceedings of the 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI'14)*, Broomfield, CO, October 2014.