Securing Cloud File Systems using Shielded Execution

Quinn Burke, Yohan Beugin, Blaine Hoak, Rachel King, Eric Pauley, Ryan Sheatsley,
Mingli Yu*, Ting He*, Thomas La Porta*, Patrick McDaniel
University of Wisconsin-Madison, The Pennsylvania State University*

Abstract

Cloud file systems offer organizations a scalable and reliable file storage solution. However, cloud file systems have become prime targets for adversaries, and traditional designs are not equipped to protect organizations against the myriad of attacks that may be initiated by a malicious cloud provider, co-tenant, or end-client. Recently proposed designs leveraging cryptographic techniques and trusted execution environments (TEEs) still force organizations to make undesirable trade-offs, consequently leading to either security, functional, or performance limitations. In this paper, we introduce `TFS`, a cloud file system that leverages the security capabilities provided by TEEs to bootstrap new security protocols that meet real-world security, functional, and performance requirements. Through extensive security and performance analyses, we show that `TFS` can ensure stronger security guarantees while still providing practical utility and performance w.r.t. state-of-the-art systems; compared to the widely-used NFS, `TFS` achieves up to 2.1× speedups across micro-benchmarks and incurs $< 1\times$ overhead for most macro-benchmark workloads. `TFS` demonstrates that organizations need not sacrifice file system security to embrace the functional and performance advantages of outsourcing.

1 Introduction

Cloud file systems are a backbone of modern cloud infrastructure. Often used as the storage interface for personal cloud drives and enterprise server applications, they provide convenient and reliable access to shared file data. While advantageous for several reasons, storing file data in the cloud raises significant security and privacy concerns [69].

Breaches of private user data and metadata, intellectual property theft, and ransomware campaigns have been shown to be particularly effective in cloud environments [4, 27, 63], highlighting the need for better ways of protecting data stored in the cloud. Further, adversaries in cloud environments include not only co-tenants and end-clients, but even a malicious cloud provider. More sophisticated defenses are required to mitigate attacks initiated by a malicious cloud provider (e.g., host system call tampering) [24, 28, 45, 56, 73]; these are commonly denoted as *host-interface attacks.* Concretely, a trusted cloud file system must therefore provide: (1) confidentiality and integrity protection for all file data and metadata, (2) resilience against a variety of host-interface attacks, (3) support for canonical features like file sharing and policy management, and (4) practical performance.

Designing a cloud file system that simultaneously meets all of these requirements is a challenging task. Widely-used cloud file systems like Amazon's EFS or Google's Filestore [9, 26, 35, 40] can deliver high-performance, but necessarily force organizations to simply trust that neither the cloud provider, nor any other privileged or unprivileged adversary, can or will maliciously access or modify file data or metadata stored on the remote hosts. And while recent efforts have leveraged cryptographic techniques and trusted execution environments (TEEs) to secure data, they still force organizations to make undesirable trade-offs and fail to deliver either sufficient security controls, feature support, or performance guarantees [6, 14, 15, 21, 30, 57, 65]. This has consequently prevented these designs from seeing wide adoption as a primary storage interface. Thus, the community lacks a suitable file system that strikes a good balance between real-world security, functional, and performance requirements.

In this paper, we introduce `TFS`, a cloud file system that meets real-world security, functional, and performance requirements. `TFS` leverages the security capabilities provided by TEEs [44, 49, 52, 61] to bootstrap new security protocols that grant four key properties: (1) confidentiality and integrity protection for all file data and metadata; (2) comprehensive protection against host-interface attacks; (3) secure and high-performance file sharing; and (4) extensible feature support. `TFS` demonstrates that organizations need not sacrifice file system security to embrace the functional and performance advantages of outsourcing.

Accomplishing this requires addressing a range of challenges associated with request processing and data persistence. First, protecting confidentiality and integrity requires designing novel end-to-end protocols that can mitigate various known attacks with minimal overhead. We address this through *data & metadata isolation*, wherein we securely partition file system tasks across trusted and untrusted components to efficiently protect against tampering with data and metadata while in-flight, in-processing, and at-rest. Second, protecting against host-interface attacks requires a careful reconsideration of the host-interface design to be able to reason about and mitigate them. We address this through *host-interface shielding*, wherein we design a simple, deterministic host-interface and develop mechanisms to protect against tampering with host-interface parameters or return codes. Lastly, providing secure and high-performance file sharing and extensible feature support requires a reliable but versatile cryptographic key management system that minimizes the risk of key compromise and has minimal performance overhead. We address this by *offloading cryptographic work* to the TEE, where the TEE serves as a trusted key escrow that manages persistent encryption keys and negotiates ephemeral keys with clients as needed.

Our evaluation of TFS examines the design trade-offs in meeting real-world security, functional, and performance requirements. We first perform a security analysis of TFS against a broad set of adversaries within the network, in-memory, and on-disk. We then provide an implementation of TFS, running in a live, cloud-like environment, and evaluate the performance across a series of Filebench-based [70] micro- and macro-benchmarks. Our analysis juxtaposes TFS against the widely used and adapted NFS [9, 26, 40]. We find that TFS satisfies real-world security requirements while providing practical performance and a comparable feature-set to NFS (with Kerberos encryption enabled); TFS achieves up to 2.1× speedups across micro-benchmarks and incurs < 1× overhead for most macro-benchmark workloads.

We contribute the following:

1. An end-to-end design and implementation of TFS; TFS provides comprehensive confidentiality and integrity protection for all file data and metadata, shields the host-interface, enables secure and high-performance file sharing, and enables extensible feature support.
2. A security analysis demonstrating the resilience of TFS against a wide range of both known and new attacks in the network, in-memory, and on-disk.
3. A performance analysis demonstrating that TFS can ensure stronger security guarantees while providing practical performance w.r.t. state-of-the-art systems.

2 Background

2.1 Cloud File Systems

Cloud file systems extend the file storage capabilities of local file systems (e.g., *ext4* [17]) to a cluster of outsourced *server* and *storage* hosts (or *nodes*) connected to *clients* by a network [1]. Here, the file system is similarly composed of both global and per-file data structures that track the file system *data* (e.g., file contents) and *metadata* (e.g., file attributes and data locations). Server and storage nodes cooperate in organizing, storing, and retrieving data and metadata for clients under a shared file system; the storage nodes may be local (directly-attached) or remote (connected via a storage-area network or other network transport [16, 47]). To clients, the distributed nature of the file system is transparent; once mounted, the files presented under the mount point have the same access semantics as files stored on any local file system. Widely supported implementations of these principles include the Network File System (NFS) [40], Amazon's Elastic File System (EFS) [9], and Google's Filestore [26].

Conventional architectures typically follow a centralized client-server model [35, 40, 67]. Here, clients issue file I/O requests on behalf of end-users or applications (whether executing on-premises or outsourced themselves) to a centralized server across a network; the server itself exposes a POSIX-like file interface for clients to access files under a shared namespace. In executing file operations, the server organizes the file data and metadata as fixed-sized *blocks* across the storage nodes; the storage nodes expose a simple interface for the server to store and retrieve blocks (typically 4 KB in size). Clients typically coordinate these tasks with server and storage nodes through *remote-procedure call* (RPC) request and response messages.

2.2 Trusted Execution Environments

Trusted execution environments (TEEs) are hardware-based security primitives that isolate execution of mutually distrusting software components running on a shared host. The software components may be other tenants' user-level applications, a hypervisor, or other system software. TEEs also provide attestation capabilities, allowing remote clients to ensure the legitimacy of the code running on an endpoint with whom they are communicating.

TEEs accomplish this through *access-mediation*, hardware-based complete mediation over designated protected (inside the TEE) and unprotected (outside the TEE) regions of physical memory, or additional *CPU modes*, processor modes that restrict the scope of operations that particular software components may perform within their execution context [49]. As a result, code and data residing in the TEE is granted strong confidentiality and integrity protection even in the presence of malicious software or hardware external to the TEE. Mature TEE implementations offering these capabilities include *Intel SGX* [49], *AMD SEV* [44], and *ARM TrustZone* [52].

3 Security Model

System Components. We assume a centralized client-server model (see Fig. 1) [29, 40]. The file system is orchestrated by five components: *client*, *network*, *server*, *TEE*, and *storage*. On the frontend, the client software provides a file interface to either end-users (e.g., employees in an enterprise network) or applications (e.g., a company's web servers). The client communicates over the network to a TEE running on an outsourced server. The tasks at the server are handled by code running either inside of the TEE or outside—denoted hereafter as the "TFS server" and "untrusted host", respectively. The storage backend consists of the outsourced local or remote storage nodes that receive commands to store or retrieve data in fixed-sized blocks. In executing file I/O requests, messages between the clients, TFS server, and storage nodes are proxied by the untrusted host.

---

[1] This architecture falls under the umbrella of distributed file systems.

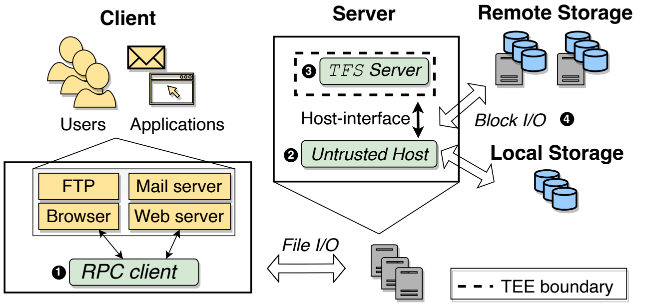

Figure 1. System components and workflow. ❶ Clients perform file I/O by having the ❷ untrusted host proxy RPC messages to the ❸ TFS server and ❹ storage nodes.

Trust Model. We consider an unmanaged deployment model, where the organization deploys and administers the file system. However, our design principles also extend to fully-managed deployments, where the cloud provider offers file storage as-a-service. We envision TFS as a replacement for widely-used systems in either case [9, 26, 40]. We therefore consider a client and TEE trusted components, and the network, untrusted host, and storage nodes untrusted. We assume the client trusts the TEE implementation.

Threat Model. Our threat model is rooted in three key observations: both file data *and* metadata have become high-yield targets for adversaries [18, 25, 37, 48, 63], host-interface attacks are a significant threat to TEE-based software [22, 45, 73], and weak or complex cryptographic key management increases the risks associated with key compromise [7, 33]. As such, the system is subject to attempts to maliciously *access*, *corrupt*, *swap*, *replay*, *reorder*, or *drop* data sent between the clients, untrusted host, TFS server, and storage nodes [6, 21, 42, 57]. The untrusted host may abuse the host-interface—for example, by crafting malicious arguments or return values to hijack control-flow between client/storage and the TEE. And lastly, adversaries may attempt to steal the keys used to encrypt data on-disk.

In line with prior work, we consider denial-of-service, physical, and side-channel attacks out of scope (e.g., network-traffic analysis [32] and other TEE-based side-channel attacks [21]). We further discuss the limitations of extant TEEs in Section 9. Our threat model resembles those of recent TEE-based file systems, but differs in the wider range of attacks that we aim to address together—notably, swapping attacks, host-interface attacks, and key compromise.

Security Requirements. To meet real-world security requirements, the file system must therefore provide end-to-end confidentiality and integrity protection for both file data and metadata, protection against host-interface attacks, and a reliable cryptographic key management system that minimizes the risks associated with key compromise.

4 Design Challenges

At surface-level, designing a file system that meets our security requirements may appear a trivial task: encrypt data, sanitize inputs, etc. However, designing an end-to-end solution is a much more nuanced endeavor. For example, while encrypting data suffices to protect confidentiality, there are security and performance trade-offs in deciding who has access to encryption keys and where encryption occurs. We characterize the key challenges under three themes.

C1. Protecting confidentiality and integrity. Ensuring confidentiality requires new mechanisms that isolate all file data and metadata from untrusted components while in-flight, in-processing, and at-rest. Ensuring integrity requires being able to attest the authenticity and correctness of code and data while processing client requests. The central challenge here lies in deciding how to securely partition tasks across trusted and untrusted components. In particular, at the server, the trust and privilege levels of components need to be considered at a far more granular level than in conventional designs [21, 40]. Current TEE-based file systems still leave open several avenues for attack (e.g., expose metadata), and a simple port of a file server like NFS to a TEE runtime still leaves many security issues unresolved (e.g., key management). We must therefore develop a new set of end-to-end protocols that enable us to more sensibly reason about and mitigate attacks, with minimal overhead.

C2. Protecting the host-interface. Mitigating host-interface attacks is a central challenge for cloud software [45, 73]. In our context, a malicious host or storage node may craft malicious arguments or return values to divert control-flow or cause other confidentiality and integrity violations. For example, valid, encrypted block data may be unknowingly swapped in place of that actually requested, before being delivered to the TEE from storage. Data encryption alone cannot defend against such attacks. Further, reasoning about and mitigating them across large and complex host-interfaces has been shown to be infeasible [45, 73]; prior efforts provide support only for a limited set of defenses [66]. The typical TEE-based library operating system (libOS) model [21, 65] is therefore ill-fit for use here. To comprehensively protect against them therefore requires judicious host-interface design and techniques that consider how inputs from the host may affect higher-level file system semantics.

C3. Supporting diverse file system features securely and efficiently. Cloud file systems are expected to support typical features like file sharing and high-level policy management [3, 7, 19, 64, 72, 77]. While various cryptographic techniques have been proposed to realize this, such approaches have significant practical limitations. For example, the typical, client-centric encrypt-then-upload model requires clients to support ad hoc cryptographic protocols

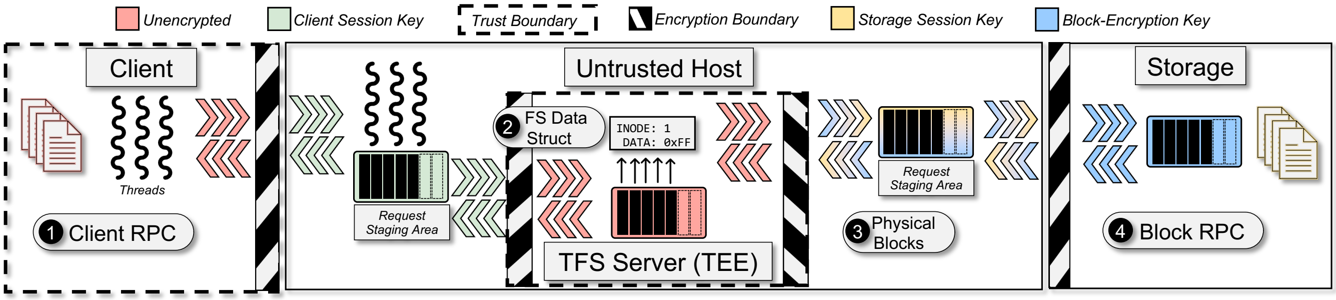

Figure 2. TFS design. Clients request ❶ file-level I/O by communicating with the TFS server through an RPC interface. The TFS server handles the client requests by ❷ updating the file system data structures within the TEE. The TFS server coordinates with the untrusted host to store and retrieve the underlying ❸ blocks on storage nodes through a ❹ block RPC interface.

and manage additional secrets. This increases the risks associated with key compromise (from lack of expertise, social engineering, or other human oversight). It complicates the semantics of file sharing; supporting a common application service like collaborative document editing is infeasible here. Moreover, it introduces performance limitations and additional constraints on feature support (e.g., supporting compliance auditing for an enterprise). Reconciling these concerns therefore requires a key management system that is reliable, versatile, and low-overhead.

5 TFS Design

TEEs provide a unique opportunity to challenge the basic premise of prior cloud file system designs. However, while TEEs provide primitives to isolate and mediate access to sensitive data in memory, extending those guarantees beyond system memory to remote clients and persistent storage media is non-trivial. The challenges stem from the fact that TEEs are sandboxed environments and rely on the untrusted host (or other kind of supervisor) to proxy access to external resources like network cards. Some information must therefore be exposed to the host such that it correctly execute requests on behalf of the TEE. How to enable this capability securely and efficiently remains an open question.

Our central goal is therefore to seek out new abstractions that provide a more practical set of trade-offs. Our design is guided by three design principles:

- **Isolate data and metadata.** We use the strong security guarantees of TEE hardware to bootstrap new security protocols that protect the confidentiality and integrity of all file data and metadata while in-flight, in-processing, and at-rest.
- **Provide shielding support.** We pivot on our isolation protocols to develop a comprehensive set of mitigations against host-interface attacks.
- **Offload cryptographic work.** We introduce an escrow-based key management system that leverages TEE capabilities to reduce the risks associated with key compromise and streamline feature support.

5.1 Isolating Data and Metadata

Conceptually, isolating data and metadata requires two tasks: deciding where file operations should execute and what the host-interface should look like. This is challenging for several reasons. First, both data and metadata are sensitive information, as they directly (through file contents, permissions, etc.) disclose private information about users and who they communicate with. They must therefore exist in plaintext only within the TEE (or client memory). Code running inside the TEE must then be able to understand the notions of directories and files to some extent, and code running outside should not be able to learn what the sensitive data is.

Second, guaranteeing the integrity of file I/O requests requires that the core file system logic (file operation handlers) be attestable by clients. Using a libOS or other POSIX wrapper library that deserializes client requests but then redirects them onto a local file system managed by the untrusted host precludes clients from being able to have assurance over how the file operation is actually implemented underneath.

Third, the decision of how to partition tasks as above directly impacts the granularity of the resulting host-interface. Opting for a libOS or wrapper library may reduce development efforts in porting core file system code to run within the TEE [6, 21, 65]), but comes at the expense of an enlarged host-interface that then needs protection. Current defense efforts for libOSes provide support only for a limited set of attacks [66]. Such approaches also observe significant performance overheads, often > 10× (and sometimes > 100×) end-to-end for local and remote clients [2, 6, 21, 55].

Toward this, we introduce three abstractions: a *trusted file system core*, *secure I/O channels*, and a *partitioned block layer*.

5.1.1 Trusted File System Core.

In TFS, the file operation handlers execute entirely within the TEE. As shown in Fig. 2, the TFS server first consumes a buffered file or block RPC message from a queue located in unprotected memory, decrypts and deserializes it, then dispatches it to the appropriate file operation handler. Any outbound file or block RPC messages are then serialized, encrypted, and submitted through a similar queue in unprotected memory. Note that the file system has a metadata layout akin to Unix-based local file systems [17], with a superblock, inode table, etc. Any data or metadata resident outside of the TEE is opaque to the untrusted host. And our design therefore reduces the host-interface size to only four functions: sending and receiving file and block RPC messages.

5.1.2 Secure I/O Channels. Bridging the clients on the frontend to the storage nodes on the backend then requires a secure transport layer. While standardized protocols like TLS provide means to realize this, the question here is what data can or should reside at the transport layer and above it.

We first distinguish between two distinct types of communication channels: I/O channels and RPC channels. As shown in Fig. 3, I/O channels form logical connections between two endpoints. In contrast, RPC channels serve as the transport for I/O channels. I/O channels thus may contain sensitive data (file names, contents, R/W offsets, etc.) that must be kept secret from untrusted components, and we therefore require them to be terminated in the TEE. While RPC channels contain non-sensitive data (assuming an encrypted payload) that need not be kept secret, terminating the RPC channels in the untrusted host (which has been the de facto best practice) introduces vulnerabilities to host-interface attacks. We similarly require RPC channels to be terminated in the TEE; we defer further discussion on this to Section 5.2.

File I/O requests from clients are therefore protected by encrypting and authenticating all I/O parameters under ephemeral key $\mathcal{K}_C$ before issuing them to the TFS server. $\mathcal{K}_C$ is known only to them. MACs are computed over the request buffer and sequence numbers tracked by the client and TFS server.

Block addresses (device ID/block ID pairs) must be exposed to the untrusted host and storage nodes such that they can correctly route and execute block I/O requests. We therefore treat plaintext block data as sensitive, but block addresses as non-sensitive. As detailed below, we encrypt block data prior to being marshalled into I/O requests. However, here block addresses may equally be stored in plaintext inside or outside the TEE. As an additional layer of integrity protection against network adversaries, block I/O requests are similarly encrypted and authenticated by the TFS server and storage nodes under ephemeral key $\mathcal{K}_S$.

5.1.3 Partitioned Block Layer. The block layer is the exit point in the TEE where data must be prepared to be stored persistently on disk. Blocks are first encrypted and authenticated by the TFS server under a persistent block-encryption key $\mathcal{K}_T$; blocks are similarly decrypted in the TEE when retrieved from storage. The key is known only to the TFS server, and therefore the block I/O channel is terminated only at the TFS server. After encryption, blocks are marshalled into (and unmarshalled from) block I/O requests by the TFS server and delivered to storage nodes by the untrusted host. We note that blocks may therefore be doubly-encrypted and authenticated: first as blocks (under $\mathcal{K}_T$), then as block RPC payloads (under $\mathcal{K}_S$). As an additional layer of protection, the block address is similarly authenticated by the TFS server.

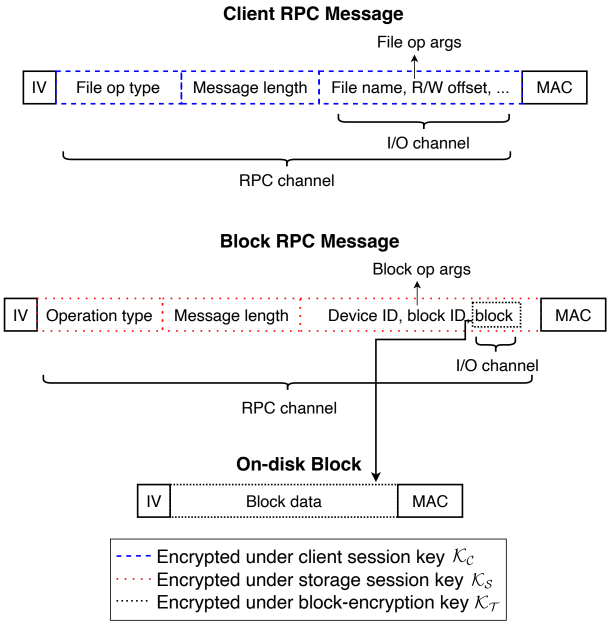

Figure 3. This message format captures the intuition behind our secure protocols for isolating file data and metadata.

5.1.4 Balancing Security and Performance. LibOSes have been central to TEE-based software development, but they are not a one-size-fits-all tool.

Strong Isolation. In TFS, clients are presented a canonical POSIX file interface. We take a microkernel approach to the server design, providing a file-system-as-a-service that is attestable to clients and ensures the confidentiality, integrity, and freshness of client data and all code handling the data.

Cutting Costs. Yet, the significance of this design extends beyond simply that we protect metadata and prescribe a smaller host-interface. It enables us to more efficiently design integrity protection mechanisms. We implement blanket integrity protection for all files at the block layer rather than the file system layer (i.e., provide full-disk encryption capabilities without the downsides of current FDE methods). This enables us to avoid having to use ad hoc solutions for ensuring integrity—e.g., per-file hashes, which can be difficult to translate to block-level representations suitable for storage on disk [30]. It also offers performance advantages. It eliminates extraneous abstractions on the critical path to storage—like syscall interfaces, VFS layers, etc. Further, it

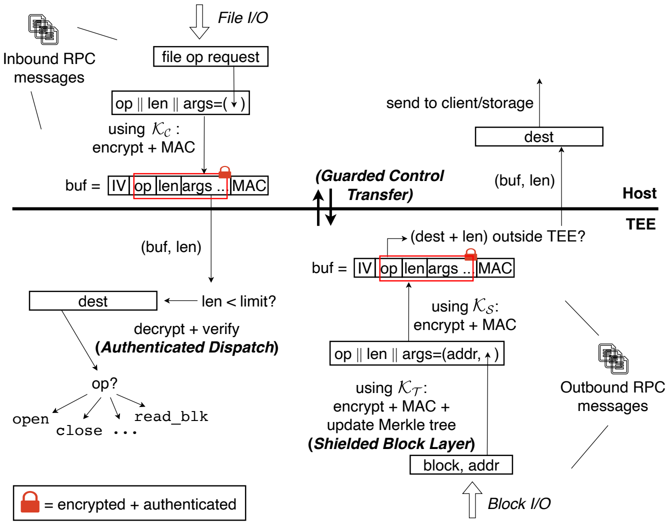

Figure 4. By authenticating RPC opcodes, using a tailored MAC construction and Merkle tree, and guarding control transfer, TFS shields against host-interface attacks.

avoids having to recompute costly checksums/hashes over large files for trivial changes (e.g., single-block updates) [79].

5.2 Providing Shielding Support

Prescribing a smaller host-interface is critical to being able to more easily mitigate host-interface attacks; we contrast this with libOS approaches that expose tens or hundreds of host-interface methods. In TFS, state transitions at the host-interface are deterministic and predictable for all four message types. Unlike prior works, we can therefore exhaustively reason about how a malicious host may tamper with the interface parameters and return codes. We introduce three additional abstractions: *authenticated dispatch*, *shielded block layer*, and *guarded control transfer.*

5.2.1 Authenticated Dispatch. The entry point for client requests at the server is the RPC layer. While RPC systems have been well-studied, how to properly terminate an RPC channel in a TEE is an open question. Terminating RPC channels in the untrusted host (by simply hooking the functions running in the TEE to appropriate RPC handler stubs) has been key to accelerating I/O in TEE-based systems [21, 36, 55]. However, this approach directly exposes RPC opcodes to the untrusted host, and are therefore vulnerable to the untrusted host simply changing the opcodes to invoke arbitrary RPC handlers. For read-only interfaces, this can cause incorrect data to be returned to users or applications, and for read-write interfaces, this can cause mutations to the file system state to be incorrect.

The root of the problem stems from RPC interfaces containing handler functions with similar or identical function signatures. Consider a host-interface with methods for opening files and changing file permissions. An open operation has the signature `int open(const char *pathname, int flags)`, and a chmod operation has the signature `int chmod(const char *pathname, mode_t mode)`, with `mode_t` defined as the same integer type. With identical signatures, a malicious host can recast an open operation into a chmod operation, and the TEE will interpet the same (valid) I/O parameters in a different context. This would allow the host to induce a permissions change on a file.

In TFS, we therefore consider RPC opcodes sensitive and terminate file RPC channels (in addition to I/O channels) in the TEE. All file RPC parameters are authenticated and verified by the TFS server and clients before any file operation proceeds. As shown in Fig. 4, once clients attest the TEE, this ensures controlled dispatch of file I/O requests: only the file operation requested by the client is invoked by the TFS server. Note that we could alternatively delegate RPC tasks to the untrusted host by exposing, but still authenticating, the type code. However, we aim to limit the number of entry points into the TEE—keeping the interface size small, constant (w.r.t. the number of supported file operations), and deterministic. We also aim to reduce costs associated with crossing protection boundaries to perform integrity checks.

5.2.2 Shielded Block Layer. The exit point for requests on the backend is the block layer. In contrast to client RPC channels, block RPC parameters are considered non-sensitive because we require the untrusted host and storage nodes to handle persistence of blocks. We do not enforce similar restrictions on block RPC messages. However, the TEE stills need to be able to detect and respond to a malicious host that supplies corrupt, replayed, or swapped blocks. Towards this, we extend the calculation for block MACs to use the *block address* as additional authenticated data (AAD) along with the block data. We then use a Merkle hash tree [50], with the block MACs as the leaves, to prevent block replays/rollbacks and ensure the freshness of block data when retrieved by the TEE. The tree is stored on a separate region of the disk and read into protected memory on boot.

Like other file systems, our Merkle tree protects block correctness and freshness. However, the Merkle tree alone is still prone to second-preimage attacks and cannot prevent valid, encrypted blocks from being swapped in place of those actually requested before being delivered to the TEE from storage. In contrast to prior designs, our MAC construction therefore additionally prevents block swapping attacks.

5.2.3 Guarded Control Transfer. Our authenticated dispatch and shielded block layer mechanisms mitigate confidentiality and integrity violations resulting from maliciously crafted host-interface parameters. It remains to consider how to mitigate attacks resulting from malicious return codes supplied to the TEE by the host. While prior work has studied similar host-interface attacks to some extent [45, 73], current mitigations address only a few specific attacks [66].

The root of the problem lies in how return codes are typically handled in file systems. Standard practice in Linux for system calls like `read()` is to propagate return codes (both successes and errors) up the call stack from device drivers, through the block layer and file system layer, and back to clients [38]. Note that the different software layers all use standard Unix `errno` codes. However, simply permitting the untrusted host to propagate arbitrary return codes would enable it to exploit vulnerabilities or weaknesses in the error-handling or decision-making logic in the file system or application code. We therefore need a mechanism to more rationally handle return codes.

The `TFS` server intercepts return codes from the untrusted host and handles them in one of two ways. A return code of zero indicates a success, and the server proceeds. Any logical failures will be detected by the trusted file system core and either handled locally (e.g., retry block-write) or reported back to the client. Any other return code indicates a failure and is transformed into a generic I/O failure (EIO) before being reported back up the call chain to the client. False positive return codes will therefore be detected on a subsequent read/write via the Merkle tree. False negative return codes may only cause retries at the server (up to some limit) or generic I/O errors at the client; indeed physical I/O errors are typically handled transparently by the cloud provider [8, 11]. In the absence of formally-verified file system or application code, this provides a hardened file system that raises the bar for attackers looking for control-flow exploits.

5.3 Offloading Cryptographic Work

Supporting a diverse set of features securely and efficiently has been a central challenge in secure file system design. Several decades of research have broadly focused on client-side encryption techniques (i.e., encrypt-then-upload) as a means for protecting outsourced file data [15, 30, 42]. While useful in narrow contexts, such designs are ill-fit for typical usage patterns of cloud storage. We take a different approach in `TFS` by offloading as much cryptographic work to the `TFS` server as possible. We introduce two abstractions: a *trusted key escrow* and a *shared persistent key*.

5.3.1 Trusted Key Escrow & Shared Persistent Key. In `TFS`, we recognize the `TFS` server as an extension of each client running on the outsourced server and appoint it as a trusted key escrow for clients. We now revisit the use of the block-encryption key ($\mathcal{K}_{\mathcal{T}}$) and ephemeral client ($\mathcal{K}_C$) and storage ($\mathcal{K}_S$) keys; our security protocols are shown in Fig. 8. We first distinguish between the notions of encryption for *persistence* and for *transport*: data is encrypted for persistence for the purpose of being stored on disk and encrypted for transport for the purpose of being sent in RPC messages.

The `TFS` server encrypts blocks for persistence under key $\mathcal{K}_{\mathcal{T}}$, shared by all clients. We note that $\mathcal{K}_{\mathcal{T}}$ may represent a single key or a master key from which other persistent keys are derived (but shared by all clients). $\mathcal{K}_{\mathcal{T}}$ is known only to the `TFS` server and generated when the server first boots. Blocks are encrypted for transport in RPC messages to clients under a per-client session key $\mathcal{K}_C$. The key is negotiated when a client mounts the file system. Note that block RPC parameters are treated as non-sensitive (as blocks are internally shielded), thus we do not require a similar construction for $\mathcal{K}_S$ (it may or may not be ephemeral/shared among storage nodes).

5.3.2 Key Maintenance. Using a trusted key escrow introduces additional challenges for bootstrapping the file system. For generating and storing persistent keys (like $\mathcal{K}_{\mathcal{T}}$), prior work has relied on unique sealing keys burned into the processor hardware on the server. Yet, part of the advantage in outsourcing lies in the flexibility in service placement: the `TFS` server may be migrated to a different machine due to third-party control-plane decisions, server failure, etc. Besides flexible placement, persistent keys must also be rotated occasionally to prevent attacks enabled by cryptanalysis; coupling the persistent key to the physical machine complicates this. When outsourcing, we therefore require more flexibility in how $\mathcal{K}_{\mathcal{T}}$ is generated and stored.

In `TFS`, $\mathcal{K}_{\mathcal{T}}$ is machine-independent (i.e., initialized when the file system is formatted). We then use the unique sealing key of the TEE as a key-encrypting key to persist $\mathcal{K}_{\mathcal{T}}$ on the current machine where `TFS` is running. This provides hardware-backed persistence of the block-encryption key, without requiring any additional key service (third-party or otherwise), and while retaining data availability as the `TFS` server is relocated to different physical machines. Moreover, it allows administrators to rotate persistent keys as often as necessary (without requiring a separate physical machine) and enables a seamless key transition period (by permitting incremental re-encryption of data under the new key).

5.3.3 Balancing Security, Performance, and Utility. The central challenge with supporting diverse requirements lies in how to efficiently manage encryption keys. We highlight the advantages of our approach below.

Assessing Risks. Entrusting the `TFS` server with the persistent key enables us to overcome many of the security risks associated with conventional client-side encryption approaches. Indeed, (non-TEE-based) delegated key management has become a pivotal aspect of cloud services [10, 20]; TEEs provide a unique opportunity to capitalize on both the security and performance advantages of delegated key management. Notably, clients are not required to have expertise and infrastructure (e.g., trusted hardware modules) to properly protect keys and other secrets from being compromised. This reduces the risk of key compromise due to lack of expertise, social engineering, or other human oversight.

Indeed, clients must trust that the TEE implementation will faithfully protect the key as intended. However, significant discrepancies arise in arguments about how trust assumptions change between the two approaches, as client machines are typically equipped with similar processors as the server operating the TEE, and clients must therefore *still* trust that the client processor's firmware is acting in good faith if/when handling secrets. Our approach raises the bar for attackers by delegating to the administrators the task of hardening the (server) machines carrying secrets. This opens many opportunities to improve utility and performance.

Secure and High-Performance File Sharing. The data access model in `TFS` is similar to that of NFS, where file operations are executed at the server. We contrast this to other approaches that cache whole-files at clients and largely execute file operations at the clients [30, 31]. Our approach ensures that the mechanics of encrypting data for persistence are transparent to clients. In turn, this avoids requiring clients to bootstrap costly, interactive, client-to-client protocols to perform simple tasks like sharing files.

Sharing is done using typical `chmod` or `setfacl` requests. Recipients can then begin retrieving the data, encrypted under their session key, without knowledge of the persistent key. Importantly, access is asynchronously granted to recipients, without requiring to explicitly notify them or otherwise requiring them to be online during the sharing process. We contrast this with approaches that require always-online clients for sharing to effectively occur (e.g., to distribute persistent keys) [30], which can become prohibitive as the file system grows or access rights change frequently.

Efficient Revocation. Revocation in secure file systems is notoriously challenging [7, 34], often requiring complex protocols for generating new encryption keys, re-encrypting data under the new keys, and distributing the new keys to the clients retaining access rights. In `TFS`, the separation between keys used to protect data for persistence and those for transport provides a dual benefit to file sharing. It revives canonical semantics of revocation: revocations are enforced through simple permission/ACL changes on files. This requires a single operation by the data owner, and revoked clients immediately lose access to the files.

Policy Management. Our data access model also simplifies the mechanics of other administrative tasks. Specifically, using the `TFS` server as an escrow allows us to realize a TEE-driven reference monitor, as all requests to read or write data must pass through the `TFS` server. This design point resembles traditional escrow systems, but differs in that the trust in the escrow is hardware-backed, and the `TFS` server can perform complex file system operations rather than simply key storage. The escrow therefore has three unique capabilities.

First, it can enforce access controls on data for both normal users *and* administrators. For example, the `TFS` server can allow administrators to perform compliance auditing, while user's can attest (through attestation over the server code) that the programmed policies meet reasonable expectations of user privacy (against both administrators and other users). Second, giving data visibility to the `TFS` server enables implementing tailored optimizations server-side, such as block-level replication, prefetching, etc. Lastly, retaining a canonical POSIX API for clients decouples client and server interface dependencies—which would otherwise make it intractable to patch or implement new features server-side without incurring compatibility hazards. To support this, the `TFS` server exposes RPC methods for configuring file-level access controls (ACLs) and other system-wide policies.

6 Security Analysis

Below we provide an analysis of the security guarantees provided by `TFS`, with a particular focus on confidentiality and integrity. We organize the analysis around the primary `TFS` components—examining a concrete set of attacks against the client, untrusted host, `TFS` server, and storage nodes. Attacks reflect those enumerated in our threat model (see Section 3).

Client. While client machines are considered trusted, an attacker who successfully compromises a client machine may obtain access to any sensitive data the client has cached locally, as well as the client's session key (and thus can temporarily impersonate them). While such is the case for any file system, clients in `TFS` do not manage persistent secrets and therefore the attacker would not have unfettered access to file system data. We can therefore minimize the blast radius in the event of a compromise.

Untrusted Host. At the untrusted host, malicious third-party software/firmware, or a co-located tenant who has gained escalated privilege, may attempt to access or corrupt messages or return codes delivered to the latter three components. Our trusted file system core ensures that sensitive data/metadata exists in plaintext only within the TEE, while encrypted RPC messages and blocks stored outside of the TEE are opaque to the untrusted host. Our authenticated dispatch, shielded block layer, and guarded control transfer mechanisms ensure that RPC channels cannot be hijacked or replayed, and return codes cannot be manipulated to arbitrarily direct control-flow. The primary file system secret (block-encryption key) is only known to the TEE. These mechanisms ensure the untrusted host cannot tamper with file system code or data while processing client requests.

`TFS` Server. While the `TFS` server is considered trusted, in the absence of formally-verified file system code, an attacker who manages to find and exploit a weakness in the `TFS` server code will have access to the block-encryption key and all file system data. However, the `TFS` code must be attested by clients and therefore any deviations from a trusted state (i.e., how routines are implemented or what secrets are present) will be detected by clients. We can therefore ensure that a compromise is localized to the exploitable code, and an adversary cannot arbitrarily change the TEE functionality.

Storage. Storage nodes may similarly become compromised and attempt to access or tamper with blocks as they are retrieved from or written to disk. However, attacks manifesting at storage nodes are recognized and handled by the `TFS` server as an attack by the untrusted host; our shielding mechanisms will ensure that block data read from disk is consistent with the Merkle tree, and any tampering on writes will be detected on subsequent reads.

7 Implementation

We implemented `TFS` for Linux hosts in ∼22 k lines of C++. It has a metadata layout similar to local `ext` file systems and is composed of client, server, and storage nodes.

Client. End-users or applications mount the file system to a local directory through the FUSE [74] API, with file operations sent as RPC messages to the `TFS` server. Our client implementation supports a rich set of file operations: `getattr`, `mkdir`, `unlink`, `rmdir`, `rename`, `chmod`, `open`, `read`, `write`, `release`, `fsync`, `opendir`, `readdir`, `releasedir`, and `create`.

`TFS` Server/Untrusted Host. The `TFS` server is an ext4 implementation ported to Intel SGX. It executes file operation handlers and block management tasks like block allocation across storage nodes; our implementation supports linear and striped allocation. Further, it supports in-order, synchronous request streams between clients and storage nodes. The server is multi-threaded, with one worker thread per client.

The untrusted host is an RPC server and client that communicates with clients and storage nodes on behalf of the `TFS` server. We implemented a lightweight RPC library for communication between the clients, `TFS` server, untrusted host, and storage nodes.

Storage. Storage nodes are simple block devices that receive block RPC requests over the network or locally. Requests are executed by memory-mapping the associated block-device file into unprotected memory (with an `mmap` syscall) and reading/writing at the appropriate block offset.

Authentication and Access Control. Our design primarily address access control at system-level as opposed to file-level (i.e., ensuring only clients and the `TFS` server may access *any* data in the file system). We implement file-level access control semantics similar to NFSv4 with `AUTH_SYS`-style authentication (i.e., file read/write/execute permissions enforced on unique user IDs associated with each connected/authenticated client), but note that the `TFS` server exposes hooks for configuring file-level and system-wide policies. We leave future work to integrating the system with mature authentication and authorization protocols such as Kerberos or OAuth.

Encryption and Integrity. Communication between the clients, untrusted host, `TFS` server, and storage nodes is protected via standard AES-128 symmetric encryption. We use Galois/Counter Mode (GCM) as it protects integrity with the MAC generated as part of the encryption process. The `TFS` server uses SGX cryptographic libraries while non-TEE `TFS` code uses libgcrypt. While we use pre-shared keys in our implementation for simplicity, keys could be acquired through a PKI or other key-negotiation protocols [2]. Finally, as an optimization, complete mediation of encryption and integrity checking across the host-interface is ensured through the type system: functions that traverse the host-interface only accept secure types, and sensitive data must be encapsulated in these types through encryption/MAC.

8 Performance Evaluation

We evaluate the performance of `TFS` under a set of realistic end-user-based workloads [2, 30]. This deployment scenario is emblematic of personal cloud storage drives or distributed file systems in enterprise networks. We focus our evaluation in two dimensions: throughput/latency and scalability across multiple clients. We seek to answer the following questions:

1. *Where are the fundamental performance bottlenecks?*
2. *How does the performance vary across workloads?*
3. *How well does* `TFS` *scale with multiple clients?*

8.1 Experimental Setup

Testbed. We run all experiments on a local cluster consisting of up to five client nodes, a single server, and four local storage nodes. Similar to other works [65], clients/server run on Intel Core i7-10710U 1.10 GHz processors with 12 logical cores and 32 GB memory. Storage nodes are backed by SSD storage. All machines are Debian-based and connected in a local network over 1 GbE interfaces. We compile the SGX code in *hardware mode* with SGX SDK v2.16 and Linux SGX kernel driver v2.14; the non-TEE code uses libgcrypt v1.8.5.

System Configurations. We compare the performance of `TFS` against several other systems. First, we compare against a non-SGX version of `TFS` (with SGX ocalls and ecalls simply replaced by direct function calls). Next, we compare against NFS (with/without Kerberos network encryption) for its prevalence and high-performance [9, 26, 40]. We also compare against NFS-Ganesha [51], a widely-used user-space NFS server implementation that offers performance improvements over the default kernel server. Note that while `TFS` supports remote storage, NFS is generally local-storage based, and therefore we focus our experiments on local storage setups. Finally, we compare against Gramine's (formerly Graphene-SGX) local file system implementation [21, 59]

---

[2] We chose not to use standardized security protocol suites like TLS [60] to allow us to experiment with a wide range of constructions, optimizations, and security policies in current and future work.

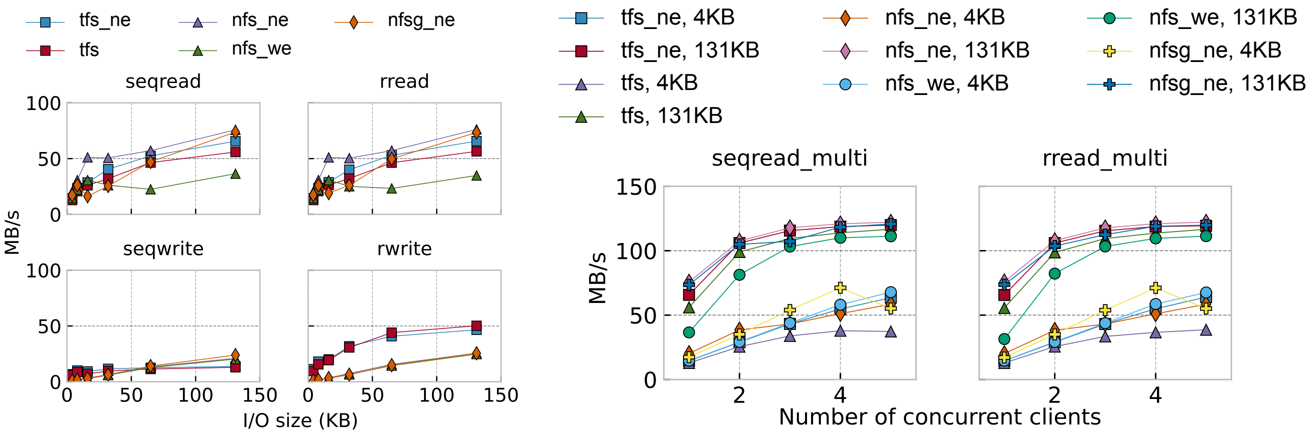

(a) Single-client, local storage

(b) Multi-client, local storage

Figure 5. Raw throughput of various read and write workloads.

for juxtaposition with another TEE-based system; an NFS port to Gramine is not available to measure end-to-end client performance. We focus on Gramine because it is under active development and widely recognized as the state-of-the-art.

The studied file system configurations provide different confidentiality and integrity guarantees and are denoted as follows: (1) `nfs_ne`, the NFS baseline without any encryption; (2) `nfs_we`, NFS with Kerberos-based network encryption and integrity protection (i.e., mounted with *sec=krb5p*) but no disk encryption; (3) `nfsg_ne`, the NFS-Ganesha baseline without any encryption; (4) `tfs_ne`, the `TFS` baseline without SGX memory encryption; and (5) `tfs`, `TFS` with full encryption and integrity protection (network, memory, and disk). Importantly, we note that `tfs` provides additional security guarantees (memory and disk encryption) over `nfs_we`, the NFS deployment mode typically used in practice.

For fair comparison, we disable client-side caching in NFS; other system parameters remain the same. Various factors can alter end-to-end performance, but our focus is to study the underlying performance of the file system infrastructure itself. Further, we only consider concurrent-readers in multi-client scenarios (as NFS does not provide consistency guarantees for concurrent-writers [40]).

8.2 Micro-benchmarks

We first study the performance of `TFS` under standard micro-benchmarks [2, 30, 43]. We aim to understand the overhead costs, scalability across the number of clients, and the raw achievable read/write throughputs of `TFS` through sequential/random read/write workloads, at various I/O sizes. In line with prior work, sequential-write workloads are block-aligned and write-allocating (i.e., require allocating new blocks on writes), whereas random-write workloads pre-allocate blocks. Mean throughput is measured across 10 trials for each micro-benchmark (Fig. 5).

Results. As shown in Fig. 5a and Fig. 5b, the general trend across all workloads and file system types is increased throughput with increased I/O size. This indicates that the client makes more efficient use of the link bandwidth per-request by using larger I/Os (i.e., RPC messaging costs are amortized). Note that we show 4 KB and 131 KB for multi-client systems to demonstrate two extremes of I/O sizes.

As `tfs_ne` runs without SGX enabled (the same code, with direct function calls used in place of SGX call gates), it serves as our baseline for understanding the relative overhead of using the TEE. For single-client systems, relative to `tfs_ne`, across all I/O sizes, `tfs` achieves > 85% of the throughput of `tfs_ne`. In fact, for write workloads, `tfs` nearly matches the performance of `tfs_ne`. For multi-client systems, we observe similar behavior w.r.t. `tfs_ne`.

For single-client systems, `tfs` in fact exceeds the performance of `nfs_we`; up to a 1.5× speedup on read-based workloads and 2.1× speedup on write-based workloads. Across read workloads, `tfs` still achieves > 75% of the throughput of either `nfs_ne` or `nfsg_ne`. For random-write workloads, `tfs` performance exceeds both, while for sequential-write workloads `tfs` achieves > 55% of the performance of either. For multi-client systems, we observe qualitatively similar behavior w.r.t. `nfs_ne`, `nfs_we`, and `nfsg_ne`.

➤ **Takeaway:** These results lead us to two conclusions. 1) SGX costs w.r.t. a non-SGX baseline are not exacerbated across workload types or I/O sizes. We reason that SGX costs (memory encryption, etc.) alone are therefore not a fundamental performance bottleneck. 2) `TFS` can deliver good raw read/write performance without sacrificing security guarantees; it exceeds performance of NFS with Kerberos encryption. We attribute this to our unified, user-space server design, whereas the NFS server runs in the kernel and often

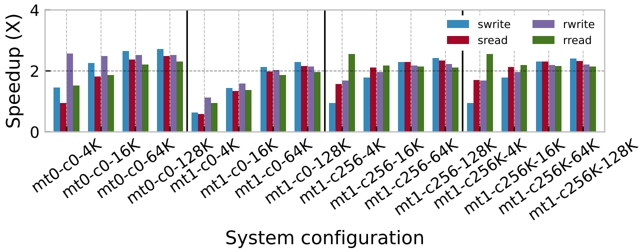

Figure 6. Speedup of TFS over Gramine. Configuration parameters are encoded in the form *a*-*b*-*c*, where *a* denotes whether the Merkle tree is enabled (1) or not (0), *b* denotes what cache size is used, and *c* denotes the I/O size used.

requires expensive upcalls to user-space daemons [40]. Further, this demonstrates that, despite conventional wisdom, pushing more logic out to the untrusted host (e.g., with a libOS runtime and exitless I/O strategy [55]) is not the only way to achieve practical performance.

Comparison against Gramine. Next, we run seq/rand read/write benchmarks directly against the TFS and Gramine file system implementations. Our goal here is to provide a quantitative performance comparison of our isolation and shielding mechanisms w.r.t. the state-of-the-art. In Fig. 6, we experiment with several different configurations, considering various workload types, whether or not TFS uses a Merkle tree, whether or not it uses a cache, and what I/O sizes are used during the benchmark. We make one overarching observation: across all configurations, TFS provides approximately $1 - 2\times$ speedups over Gramine.

➤ **Takeaway:** The key takeaway is that this serves as validation that the TFS server design provides a performant foundation for a TEE-based file system. Moreover, we note that Gramine reads an entire file and verifies the Merkle hashes only once, on file open (and flushes on close), which can become prohibitive for large files. In contrast, TFS verifies/updates hashes at block-level, on every read/write operation, and is still able to provide practical performance.

8.3 Macro-benchmarks

Next, we analyze how well TFS can support client workloads that perform a variety of other file operations (listing directories, changing permissions, etc.). We focus the analysis around various popular Linux utilities [2, 30]. Task latency is measured across 10 trials for each benchmark (Fig. 7).

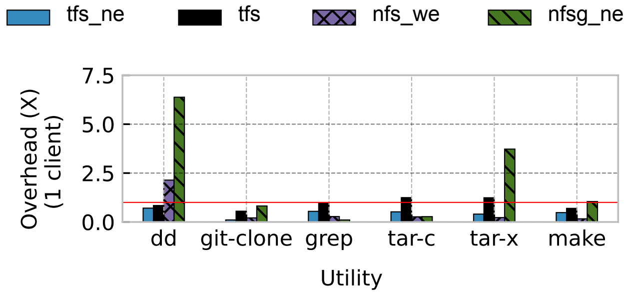

(a) Single-client, local storage

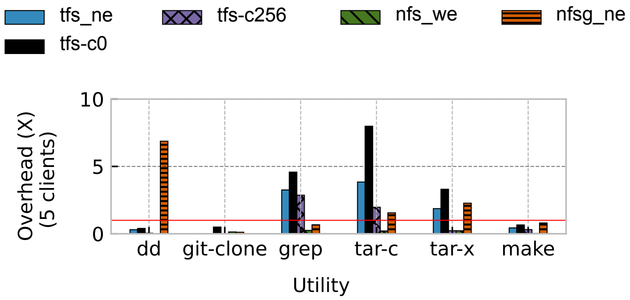

(b) Multi-client, local storage

Figure 7. Task latency of various Linux utilities. For TFS, usage of a cache is denoted by the suffix *c256*, otherwise no cache is used. A red line is drawn at 1× overhead.

Results. As shown in Fig. 7a and Fig. 7b, the general trend across all workloads is that the insecure file system types exhibit lower latencies than their secure counterparts. For single-client systems `tfs` overheads over `tfs_ne` are generally $< 1\times$. However, `tfs` in fact observes significant speedups over `nfsg_ne` for most utilities. We attribute this to the good write performance of `tfs` (as shown in the micro-benchmarks), and the fact that `tfs` can complete certain related tasks quicker (e.g., opening files, changing permissions) and, in aggregate, perform better. Our primary observation is that `tfs` only observes $< 1\times$ overheads compared to `nfs_we`. We note that caching is disabled in `tfs` here, while `nfs_we` still makes use of the entire OS page cache for not only raw disk blocks but also file handles and other file system objects. These overheads observed stand in stark contrast to those observed by prior works, often $> 10\times$ (and sometimes $> 100\times$) end-to-end for local and remote clients [2, 6, 21, 55].

For multi-client systems, we observed slightly different behavior. Relative to `tfs_ne`, `tfs` sees overhead costs of $< 1.25\times$ for all utilities (without a cache). In fact, compared to `nfs_ne`, it still sees overhead costs of only $< 1\times$ for `dd`, `git-clone`, and `make`. Other utilities show up to 7.5× overheads (for `tar-c`). However, with a small cache (256 blocks), we can largely mitigate the associated overhead costs; note that NFS makes heavy use of the OS page cache by default.

➤ **Takeaway:** Our macro-benchmark results further support the efficacy of the TFS design. We make two key observations. 1) TFS can deliver stronger security guarantees with low overhead for single- and multi-client systems w.r.t. the state-of-the-art. This demonstrates that the TFS design is sufficient for typical cloud file system deployment scenarios like personal cloud storage drives or enterprise network drives. 2) Whole-system performance is highly dependent on the number of entries/exits across the TEE boundary. Caching can help to some extent, but the root problem is that metadata operations (locating file data blocks, updating timestamps, etc.) amplify the number of entries/exits and Merkle tree operations when reading blocks at the server. This becomes more noticeable with multiple clients. We argue that moving to an alternative/decoupled metadata organization [78] will allow future designs to further minimize overheads and feasibly reach `nfs_ne`-like performance. Other server-side optimizations like readahead and replication can also be implemented in `TFS` to further improve performance.

9 Discussion

Below we discuss notable points of consideration for `TFS` deployments (see Appendix B for extended discussion).

Extending to Other TEEs. Commercial TEEs all provide some notion of isolation between a trusted and untrusted environment. For example, ARM TrustZone TEEs are characterized by a secure and non-secure "world" or state, where all memory has an extra bit defining its state [5]. AMD SEV, on the other hand, implements a VM based TEE, providing separation at the boundary of the secure VM. While our implementation is based on SGX, the design of `TFS` is not tightly tied to SGX. Indeed, the `TFS` design principles naturally extend to other TEEs by adapting the host-interface to the idiosyncrasies of the particular TEE used. We leave future work to implementing `TFS` with other TEEs.

Limitations. While our design successfully defends against a wide range of known and new attacks, side-channels attacks on TEEs still present a challenge [39, 46, 53, 54]. While the untrusted host can only see encrypted RPC messages, some underlying access patterns may be discernible and indicative of underlying workloads (e.g., a read-heavy workload from a web-server accessing static assets on the file system). We consider our architecture in the context of TEEs resistant to such attacks—towards which recent work holds promise [1, 41, 62]. We leave future work to exploring such attacks and incorporating defenses.

10 Related Work

Cloud file system design has a long history that intersects storage, applied cryptography, and trusted computing research. `TFS` builds on the lessons learned in these works, rethinking the fundamental file system abstractions to produce a design with a unique set of capabilities.

Client-based Solutions. Client-based solutions have been the traditional approach approach to protecting data stored on untrusted clouds. For example, CFS [15], Plutus [42], and NeXUS [30] require clients (or trusted client proxies) to execute file operations and handle all cryptographic tasks, while files are organized as opaque blobs to the untrusted server/storage. Other works also assume a client-based gateway to untrusted storage [23, 68, 75, 76]. While useful in narrow contexts, such designs still require clients to have proper training and expertise with managing keys. More importantly, they are ill-fit for typical usage patterns of cloud storage. For example, providing an application service such as collaborative document editing is infeasible here.

In `TFS`, we instead delegate encryption for persistence to the file server (in our design, denoted as the `TFS` server) and redesign the structure of the file server to provide stronger security guarantees (against more attacks), extensible feature support, and practical performance.

LibOS Runtimes. Many recent efforts have relied on libOS runtimes as a means for quickly porting server applications to use TEEs. LibOSes enable applications to run unmodified in TEEs by automatically generating the necessary wrapper code (including encryption/decryption operations over data) to redirect system calls from within the TEE onto the host [6, 14, 21, 57, 65, 71]. Typical storage applications include in-memory databases [58], local file systems [2], and key-value stores [12, 13]. However, simply porting a traditional file system (like NFS [40], GFS [35], or EFS [9]) to use a libOS, and equipping it with TLS and disk-encryption, would fail to meet all of our confidentiality, integrity, shielding, and key management requirements.

In `TFS`, we construct an end-to-end design from the ground up, without relying on a libOS runtime. Instead, we design a new file system core that provides protection for all data and metadata, provides comprehensive protection against host-interface attacks, and seamlessly handles encryption for persistence and transport to enable high-performance file sharing and policy management.

11 Conclusion

Cloud file systems have become a critical component of modern cloud infrastructure. As threat models evolve and security requirements become stricter, new security mechanisms are needed to protect against the myriad of attacks that may be initiated by a malicious cloud provider, co-tenant, or end-client. At the same time, the file system must still remain flexible enough to support typical features like file sharing and policy management. We introduced `TFS`, a cloud file system that meets real-world security, functional, and performance requirements. Our evaluation demonstrated that `TFS` can provide stronger security guarantees than prior designs, while still providing practical utility and performance. On balance, `TFS` challenges current wisdom in cloud file system design and demonstrates that simple architectural changes can have significant practical advantages.

References

[1] Adil Ahmad, Byunggill Joe, Yuan Xiao, Yinqian Zhang, Insik Shin, and Byoungyoung Lee. 2019. OBFUSCURO: A Commodity Obfuscation Engine on Intel SGX. In *Proceedings 2019 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA. https://doi.org/10.14722/ndss.2019.23513

[2] Adil Ahmad, Kyungtae Kim, Muhammad Ihsanulhaq Sarfaraz, and Byoungyoung Lee. 2018. OBLIVIATE: A Data Oblivious Filesystem for Intel SGX. In *Proceedings 2018 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA. https://doi.org/10.14722/ndss.2018.23284

[3] Ramnatthan Alagappan, Aishwarya Ganesan, Eric Lee, Aws Albarghouthi, Vijay Chidambaram, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. 2018. Protocol-Aware Recovery for Consensus-Based Storage. In *16th USENIX Conference on File and Storage Technologies (FAST 18)*. USENIX Association, Oakland, CA, 15–32. https://www.usenix.org/conference/fast18/presentation/alagappan

[4] Kyle Alspach. 2022. *Microsoft Azure has had a string of 'nightmare' vulnerabilities.* Protocol Media, LLC. Retrieved 2023-05-16 from https://www.protocol.com/enterprise/microsoft-azure-vulnerabilities-cloud-security

[5] arm-security 2017. *Security in ARMv8-A systems.* https://developer.arm.com/documentation/100935/0100/The-TrustZone-hardware-architecture-

[6] Sergei Arnautov, Bohdan Trach, Franz Gregor, Thomas Knauth, Andre Martin, Christian Priebe, Joshua Lind, Divya Muthukumaran, Dan O'Keeffe, Mark L. Stillwell, David Goltzsche, Dave Eyers, Rüdiger Kapitza, Peter Pietzuch, and Christof Fetzer. 2016. SCONE: Secure Linux Containers with Intel SGX. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. USENIX Association, Savannah, GA, 689–703. https://www.usenix.org/conference/osdi16/technical-sessions/presentation/arnautov

[7] Giuseppe Ateniese, Kevin Fu, Matthew Green, and Susan Hohenberger. 2006. Improved proxy re-encryption schemes with applications to secure distributed storage. *ACM Transactions on Information and System Security (TISSEC)* 9, 1 (2006), 1–30.

[8] Amazon AWS. 2023. *Amazon Elastic Block Store Service Level Agreement.* Amazon Web Services, Inc. Retrieved 2023-05-16 from https://aws.amazon.com/ebs/sla/

[9] Amazon AWS. 2023. *Amazon Elastic File System.* Amazon Web Services, Inc. Retrieved 2023-05-16 from https://aws.amazon.com/efs

[10] Amazon AWS. 2023. *AWS Key Management Service (AWS KMS).* Amazon Web Services, Inc. Retrieved 2023-05-16 from https://aws.amazon.com/kms/

[11] Amazon AWS. 2023. *Failover with AWS.* Amazon Web Services, Inc. Retrieved 2023-05-16 from https://docs.aws.amazon.com/whitepapers/latest/web-application-hosting-best-practices/failover-with-aws.html

[12] Maurice Bailleu, Dimitra Giantsidi, Vasilis Gavrielatos, Do Le Quoc, Vijay Nagarajan, and Pramod Bhatotia. 2021. Avocado: A Secure In-Memory Distributed Storage System. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*. USENIX Association, 65–79. https://www.usenix.org/conference/atc21/presentation/bailleu

[13] Maurice Bailleu, Jörg Thalheim, Pramod Bhatotia, Christof Fetzer, Michio Honda, and Kapil Vaswani. 2019. SPEICHER: Securing LSM-based Key-Value Stores using Shielded Execution. In *17th USENIX Conference on File and Storage Technologies (FAST 19)*. USENIX Association, Boston, MA, 173–190. https://www.usenix.org/conference/fast19/presentation/bailleu

[14] Andrew Baumann, Marcus Peinado, and Galen Hunt. 2015. Shielding applications from an untrusted cloud with haven. *ACM Transactions on Computer Systems (TOCS)* 33, 3 (2015), 1–26.

[15] Matt Blaze. 1993. A Cryptographic File System for UNIX. In *Proceedings of the 1st ACM Conference on Computer and Communications Security* (Fairfax, Virginia, USA) *(CCS '93)*. New York, NY, USA, 9–16. https://doi.org/10.1145/168588.168590

[16] Peter T. Breuer. 2000. The enhanced network block device. *Linux journal* (2000).

[17] Mingming Cao, Suparna Bhattacharya, and Ted Ts'o. 2007. Ext4: The Next Generation of Ext2/3 Filesystem. In *2007 Linux Storage & Filesystem Workshop (LSF 07)*. USENIX Association, San Jose, CA. https://www.usenix.org/conference/2007-linux-storage-filesystem-workshop/ext4-next-generation-ext23-filesystem

[18] David Cash, Paul Grubbs, Jason Perry, and Thomas Ristenpart. 2015. Leakage-abuse attacks against searchable encryption. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*. 668–679.

[19] Miguel Castro and Barbara Liskov. 1999. Practical Byzantine Fault Tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*. 173–186.

[20] Ramaswamy Chandramouli, Michaela Iorga, and Santosh Chokhani. 2013. Cryptographic key management issues and challenges in cloud services. *Secure Cloud Computing* (2013), 1–30.

[21] Chia che Tsai, Donald E. Porter, and Mona Vij. 2017. Graphene-SGX: A Practical Library OS for Unmodified Applications on SGX. In *2017 USENIX Annual Technical Conference (USENIX ATC 17)*. USENIX Association, Santa Clara, CA, 645–658. https://www.usenix.org/conference/atc17/technical-sessions/presentation/tsai

[22] Stephen Checkoway and Hovav Shacham. 2013. Iago Attacks: Why the System Call API is a Bad Untrusted RPC Interface. In *Proceedings of the Eighteenth International Conference on Architectural Support for Programming Languages and Operating Systems* (Houston, Texas, USA) *(ASPLOS '13)*. New York, NY, USA, 253–264. https://doi.org/10.1145/2451116.2451145

[23] Ming Chen, Erez Zadok, Arun Olappamanna Vasudevan, and Kelong Wang. 2016. SeMiNAS: A Secure Middleware for Wide-Area Network-Attached Storage. In *Proceedings of the 9th ACM International on Systems and Storage Conference*. ACM, Haifa Israel, 1–13. https://doi.org/10.1145/2928275.2928282

[24] Shuo Chen, Jun Xu, Emre Can Sezer, Prachi Gauriar, and Ravishankar K Iyer. 2005. Non-Control-Data Attacks Are Realistic Threats. In *USENIX security symposium*, Vol. 5. 146.

[25] Weikeng Chen, Thang Hoang, Jorge Guajardo, and Attila A. Yavuz. 2022. Titanium: A Metadata-Hiding File-Sharing System with Malicious Security. In *Proceedings 2022 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA, USA. https://doi.org/10.14722/ndss.2022.24161

[26] Google Cloud. 2023. *Google Filestore.* Google LLC. Retrieved 2023-05-16 from https://cloud.google.com/filestore

[27] Andrea Continella, Mario Polino, Marcello Pogliani, and Stefano Zanero. 2018. There's a hole in that bucket! a large-scale analysis of misconfigured s3 buckets. In *Proceedings of the 34th Annual Computer Security Applications Conference*. 702–711.

[28] Rongzhen Cui, Lianying Zhao, and David Lie. 2021. Emilia: Catching Iago in Legacy Code.. In *NDSS*.

[29] Benjamin Depardon, Gaël Le Mahec, and Cyril Séguin. 2013. *Analysis of Six Distributed File Systems.* Research Report. SysFera, University of Picardie Jules Verne. 44 pages. https://hal.inria.fr/hal-00789086

[30] J. B. Djoko, J. Lange, and A. J. Lee. 2019. NeXUS: Practical and Secure Access Control on Untrusted Storage Platforms using Client-Side SGX. In *2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 401–413. https://doi.org/10.1109/DSN.2019.00049 ISSN: 1530-0889.

[31] OpenAFS Foundation. 2023. *OpenAFS.* The OpenAFS Foundation, Inc. Retrieved 2023-05-16 from http://www.openafs.org/

[32] Xinwen Fu, Bryan Graham, Riccardo Bettati, and Wei Zhao. 2003. Active traffic analysis attacks and countermeasures. In *2003 International Conference on Computer Networks and Mobile Computing, 2003. ICCNMC 2003.* IEEE, 31–39.

[33] Walter Fumy and Peter Landrock. 1993. Principles of key management. *IEEE Journal on selected areas in communications* 11, 5 (1993), 785–793.

[34] William C Garrison, Adam Shull, Steven Myers, and Adam J Lee. 2016. On the practicality of cryptographically enforcing dynamic access control policies in the cloud. In *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 819–838.

[35] Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung. 2003. The Google file system. In *Proceedings of the nineteenth ACM symposium on Operating systems principles (SOSP '03)*. Association for Computing Machinery, New York, NY, USA, 29–43. https://doi.org/10.1145/945445.945450

[36] google-asylo-link 2023. *Asylo: An open and flexible framework for developing enclave applications.* Retrieved 2023-05-16 from https://github.com/google/asylo

[37] Benjamin Greschbach, Gunnar Kreitz, and Sonja Buchegger. 2012. The devil is in the metadata — New privacy challenges in Decentralised Online Social Networks. In *2012 IEEE International Conference on Pervasive Computing and Communications Workshops*. 333–339. https://doi.org/10.1109/PerComW.2012.6197506

[38] Haryadi S Gunawi, Cindy Rubio-González, Andrea C Arpaci-Dusseau, Remzi H Arpaci-Dusseau, and Ben Liblit. 2008. EIO: Error Handling is Occasionally Correct.. In *FAST*, Vol. 8. 1–16.

[39] Johannes Götzfried, Moritz Eckert, Sebastian Schinzel, and Tilo Müller. 2017. Cache Attacks on Intel SGX. In *Proceedings of the 10th European Workshop on Systems Security*. ACM, Belgrade Serbia, 1–6. https://doi.org/10.1145/3065913.3065915

[40] Thomas Haynes and David Noveck. 2015. *Network File System (NFS) Version 4 Protocol.* RFC 7530. RFC Editor. 323 pages. https://www.rfc-editor.org/pdfrfc/rfc7530.txt.pdf

[41] Gernot Heiser, Toby Murray, and Gerwin Klein. 2020. Towards provable timing-channel prevention. *ACM SIGOPS Operating Systems Review* 54, 1 (2020), 1–7.

[42] Mahesh Kallahalla, Erik Riedel, Ram Swaminathan, Qian Wang, and Kevin Fu. 2003. Plutus: Scalable Secure File Sharing on Untrusted Storage. In *2nd USENIX Conference on File and Storage Technologies (FAST 03)*. USENIX Association, San Francisco, CA. https://www.usenix.org/conference/fast-03/plutus-scalable-secure-file-sharing-untrusted-storage

[43] Sudarsun Kannan, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, Yuangang Wang, Jun Xu, and Gopinath Palani. 2018. Designing a True Direct-Access File System with DevFS. In *16th USENIX Conference on File and Storage Technologies (FAST 18)*. USENIX Association, Oakland, CA, 241–256. https://www.usenix.org/conference/fast18/presentation/kannan

[44] David Kaplan, Jeremy Powell, and Tom Woller. 2016. AMD memory encryption. *White paper* (2016).

[45] Mustakimur Rahman Khandaker, Yueqiang Cheng, Zhi Wang, and Tao Wei. 2020. COIN attacks: On insecurity of enclave untrusted interfaces in SGX. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*. 971–985.

[46] Sangho Lee, Ming-Wei Shih, Prasun Gera, Taesoo Kim, Hyesoon Kim, and Marcus Peinado. 2017. Inferring Fine-grained Control Flow Inside SGX Enclaves with Branch Shadowing. In *26th USENIX Security Symposium (USENIX Security 17)*. USENIX Association, Vancouver, BC, 557–574. https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/lee-sangho

[47] Sergey Legtchenko, Hugh Williams, Kaveh Razavi, Austin Donnelly, Richard Black, Andrew Douglas, Nathanaël Cheriere, Daniel Fryer, Kai Mast, Angela Demke Brown, et al. 2017. Understanding {Rack-Scale} Disaggregated Storage. In *9th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 17)*.

[48] Matteo Maffei, Giulio Malavolta, Manuel Reinert, and Dominique Schröder. 2015. Privacy and Access Control for Outsourced Personal Records. In *2015 IEEE Symposium on Security and Privacy*. 341–358. https://doi.org/10.1109/SP.2015.28 ISSN: 2375-1207.

[49] Frank McKeen, Ilya Alexandrovich, Alex Berenzon, Carlos V. Rozas, Hisham Shafi, Vedvyas Shanbhogue, and Uday R. Savagaonkar. 2013. Innovative instructions and software model for isolated execution. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy - HASP '13*. ACM Press, Tel-Aviv, Israel. https://doi.org/10.1145/2487726.2488368

[50] Ralph C Merkle. 2001. A certified digital signature. In *Advances in cryptology—CRYPTO'89 proceedings*. Springer, 218–238.

[51] NFS-Ganesha. 2023. *The NFS-Ganesha Project.* The NFS-Ganesha Project. Retrieved 2023-05-16 from https://nfs-ganesha.github.io/

[52] Bernard Ngabonziza, Daniel Martin, Anna Bailey, Haehyun Cho, and Sarah Martin. 2016. TrustZone Explained: Architectural Features and Use Cases. In *2016 IEEE 2nd International Conference on Collaboration and Internet Computing (CIC)*. IEEE, Pittsburgh, PA, USA, 445–451. https://doi.org/10.1109/CIC.2016.065

[53] Alexander Nilsson, Pegah Nikbakht Bideh, and Joakim Brorsson. 2020. *A Survey of Published Attacks on Intel SGX.* http://arxiv.org/abs/2006.13598 arXiv:2006.13598 [cs].

[54] Oleksii Oleksenko, Bohdan Trach, Robert Krahn, Mark Silberstein, and Christof Fetzer. 2018. Varys: Protecting SGX Enclaves from Practical Side-Channel Attacks. In *2018 USENIX Annual Technical Conference (USENIX ATC 18)*. USENIX Association, Boston, MA, 227–240. https://www.usenix.org/conference/atc18/presentation/oleksenko

[55] Meni Orenbach, Pavel Lifshits, Marina Minkin, and Mark Silberstein. 2017. Eleos: ExitLess OS Services for SGX Enclaves. In *Proceedings of the Twelfth European Conference on Computer Systems*. ACM, Belgrade Serbia, 238–253. https://doi.org/10.1145/3064176.3064219

[56] Dan RK Ports and Tal Garfinkel. 2008. Towards Application Security on Untrusted Operating Systems.. In *HotSec*.

[57] Christian Priebe, Divya Muthukumaran, Joshua Lind, Huanzhou Zhu, Shujie Cui, Vasily A. Sartakov, and Peter Pietzuch. 2020. SGX-LKL: Securing the Host OS Interface for Trusted Execution. *arXiv:1908.11143 [cs]* (Jan. 2020). http://arxiv.org/abs/1908.11143

[58] Christian Priebe, Kapil Vaswani, and Manuel Costa. 2018. EnclaveDB: A Secure Database Using SGX. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, 264–278. https://doi.org/10.1109/SP.2018.00025

[59] Gramine Project. 2023. *Gramine - a Library OS for Unmodified Applications.* Gramine Project. Retrieved 2023-05-16 from https://github.com/gramineproject/gramines

[60] Eric Rescorla. 2018. *The Transport Layer Security (TLS) Protocol Version 1.3.* Request for Comments RFC 8446. Internet Engineering Task Force. https://doi.org/10.17487/RFC8446

[61] Mohamed Sabt, Mohammed Achemlal, and Abdelmadjid Bouabdallah. 2015. Trusted Execution Environment: What It is, and What It is Not. In *2015 IEEE Trustcom/BigDataSE/ISPA*. IEEE, Helsinki, Finland, 57–64. https://doi.org/10.1109/Trustcom.2015.357

[62] Sajin Sasy, Sergey Gorbunov, and Christopher W. Fletcher. 2018. ZeroTrace : Oblivious Memory Primitives from Intel SGX. In *Proceedings 2018 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA. https://doi.org/10.14722/ndss.2018.23239

[63] Nolen Scaife, Henry Carter, Patrick Traynor, and Kevin R. B. Butler. 2016. CryptoLock (and Drop It): Stopping Ransomware Attacks on User Data. In *2016 IEEE 36th International Conference on Distributed Computing Systems (ICDCS)*. 303–312. https://doi.org/10.1109/ICDCS.2016.46 ISSN: 1063-6927.

[64] Mehul A Shah, Mary Baker, Jeffrey C Mogul, Ram Swaminathan, et al. 2007. Auditing to Keep Online Storage Services Honest.. In *HotOS*.

[65] Shweta Shinde, Dat Le Tien, Shruti Tople, and Prateek Saxena. 2017. Panoply: Low-TCB Linux Applications With SGX Enclaves.. In *NDSS*.

[66] Shweta Shinde, Shengyi Wang, Pinghai Yuan, Aquinas Hobor, Abhik Roychoudhury, and Prateek Saxena. 2020. BesFS: A POSIX Filesystem for Enclaves with a Mechanized Safety Proof. In *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, 523–540. https://www.usenix.org/conference/usenixsecurity20/presentation/shinde

[67] Konstantin Shvachko, Hairong Kuang, Sanjay Radia, and Robert Chansler. 2010. The Hadoop Distributed File System. In *2010 IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST)*. 1–10. https://doi.org/10.1109/MSST.2010.5496972 ISSN: 2160-1968.

[68] Emil Stefanov, Marten van Dijk, Ari Juels, and Alina Oprea. 2012. Iris: a scalable cloud file system with efficient integrity checks. In *Proceedings of the 28th Annual Computer Security Applications Conference on - ACSAC '12*. ACM Press, Orlando, Florida, 229. https://doi.org/10.1145/2420950.2420985

[69] Hassan Takabi, James BD Joshi, and Gail-Joon Ahn. 2010. Security and privacy challenges in cloud computing environments. *IEEE Security & Privacy* 8, 6 (2010), 24–31.

[70] Vasily Tarasov, Erez Zadok, and Spencer Shepler. 2016. Filebench: A Flexible Framework for File System Benchmarking. *login Usenix Mag.* 41 (2016).

[71] Jörg Thalheim, Harshavardhan Unnibhavi, Christian Priebe, Pramod Bhatotia, and Peter Pietzuch. 2021. rkt-io: a direct I/O stack for shielded execution. In *Proceedings of the Sixteenth European Conference on Computer Systems*. ACM, Online Event United Kingdom, 490–506. https://doi.org/10.1145/3447786.3456255

[72] Anjo Vahldiek-Oberwagner, Eslam Elnikety, Aastha Mehta, Deepak Garg, Peter Druschel, Rodrigo Rodrigues, Johannes Gehrke, and Ansley Post. 2015. Guardat: Enforcing data policies at the storage layer. In *Proceedings of the Tenth European Conference on Computer Systems*. 1–16.

[73] Jo Van Bulck, David Oswald, Eduard Marin, Abdulla Aldoseri, Flavio D Garcia, and Frank Piessens. 2019. A tale of two worlds: Assessing the vulnerability of enclave shielding runtimes. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 1741–1758.

[74] Bharath Kumar Reddy Vangoor, Vasily Tarasov, and Erez Zadok. 2017. To FUSE or Not to FUSE: Performance of User-Space File Systems. In *15th USENIX Conference on File and Storage Technologies (FAST 17)*. USENIX Association, Santa Clara, CA, 59–72. https://www.usenix.org/conference/fast17/technical-sessions/presentation/vangoor

[75] Paolo Viotti, Dan Dobre, and Marko Vukolić. 2017. Hybris: Robust hybrid cloud storage. *ACM Transactions on Storage (TOS)* 13, 3 (2017), 1–32.

[76] Michael Vrable, Stefan Savage, and Geoffrey M Voelker. 2012. BlueSky: a cloud-backed file system for the enterprise.. In *FAST*. 19.

[77] Yang Wang, Manos Kapritsos, Zuocheng Ren, Prince Mahajan, Jeevitha Kirubanandam, Lorenzo Alvisi, and Mike Dahlin. 2013. Robustness in the Salus scalable block store. In *10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13)*. 357–370.

[78] Sage A. Weil, Scott A. Brandt, Ethan L. Miller, Darrell D. E. Long, and Carlos Maltzahn. 2006. Ceph: A Scalable, High-Performance Distributed File System. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation* (Seattle, Washington) *(OSDI '06)*. USENIX Association, USA, 307–320.

[79] Yupu Zhang, Abhishek Rajimwale, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. 2010. End-to-end Data Integrity for File Systems: A ZFS Case Study. In *8th USENIX Conference on File and Storage Technologies (FAST 10)*. USENIX Association, San Jose, CA. https://www.usenix.org/conference/fast-10/end-end-data-integrity-file-systems-zfs-case-study

A TFS Security Protocols

File I/O Messaging

1. $C \rightarrow T : \{m_1\}_{\mathcal{K}_C}, MAC_{\mathcal{K}_C}(\{m_1\}_{\mathcal{K}_C}, s_C)$ *(client RPC request)*
2. $T \rightarrow C : \{m_2\}_{\mathcal{K}_C}, MAC_{\mathcal{K}_C}(\{m_2\}_{\mathcal{K}_C}, s_{T,C})$ *(client RPC response)*

Block I/O Messaging

3. $T \rightarrow S : \{b_1\}_{\mathcal{K}_S}, MAC_{\mathcal{K}_S}(\{b_1\}_{\mathcal{K}_S}, s_{T,S})$ *(block RPC request)*
   - a. $b_1 \leftarrow addr$ *(read)*
   - b. $b_1 \leftarrow addr, \{b\}_{\mathcal{K}_T}, MAC_{\mathcal{K}_T}(\{b\}_{\mathcal{K}_T}, addr)$ *(write)*
4. $S \rightarrow T : \{b_2\}_{\mathcal{K}_S}, MAC_{\mathcal{K}_S}(\{b_2\}_{\mathcal{K}_S}, s_S)$ *(block RPC response)*
   - a. $b_2 \leftarrow addr, \{b\}_{\mathcal{K}_T}, MAC_{\mathcal{K}_T}(\{b\}_{\mathcal{K}_T}, addr)$ *(read)*
   - b. $b_2 \leftarrow ACK$ *(write)*

Figure 8. TFS security protocols. C, T, and S represent the client, TFS server, and storage node. Sequence numbers are denoted by s.

B Extended Discussion

I/O Processing. Our current implementation of TFS assumes client and storage node communication occurs over synchronous request streams. Additionally, we offer the same consistency semantics as NFS, where concurrent writers to the same file results in undefined behavior. Often, clients leverage asynchronous communication to improve performance, such as NFSv4's *async* mount option that delays file writes [40]. Moreover, some applications may use concurrent writers. We leave future work to extending TFS with such optimizations without compromising security guarantees.

Other Backend Architectures. Our current TFS design focuses on a single-server scenario and centralized storage management (i.e., both data and metadata are managed at the same server). This architecture is emblematic of traditional cloud file systems (like NFS, EFS, etc.). Scaling compute power to multiple servers and decoupling data and metadata management (by designating servers as either a data or metadata node) have been shown to improve file system performance [78]. We leave future work to extending the TFS design to support such optimizations.