TECHNICAL SECURITY WHITEPAPER • SYSTEMS SECURITY RESEARCH

Immutable Execution Environment: A Minimal OS Modification to Permanently Stop Ransomware and Malware

How Three Surgical Kernel Changes and a 25-Line Whitelist Eliminate the Need for Antivirus, EDR, and Signature-Based Defenses — Forever

Classification	**Unrestricted — Public Distribution**
Document Type	Technical Whitepaper / Architecture Proposal
Version	1.0 — Initial Publication
Date	May 2026
Audience	Security Researchers, OS Engineers, CISOs, Policymakers
Keywords	Immutable OS, Execution Whitelist, Kernel Security, Ransomware Defense, Application Control, Mandatory Access Control, Safe Mode, TPM, IEE

Executive Summary

The fundamental reason ransomware, malware, and the majority of cyberattacks succeed is that modern operating systems permit arbitrary code execution by default. Every reactive defense — antivirus, endpoint detection and response (EDR), behavioral heuristics, threat intelligence — exists to compensate for this single architectural flaw. This whitepaper proposes a different approach: eliminate the flaw itself.

The proposed architecture, called the **Immutable Execution Environment (IEE)**, requires only three small, surgical modifications to an existing OS kernel — approximately 80 lines of new code — plus a short path-based whitelist of 15–25 directory entries. The result is a system where: only code in whitelisted directories can execute (including all scripts run by any interpreter); every whitelisted path is immutable and read-only (even root or Administrator cannot write to it during normal operation); critical data files can be whitelisted as permanently unencryptable; and system administration is restricted to a configurable Safe Mode whose network access — including complete TCP stack removal — is sealed in the system's trust anchor.

Because the defense is *structural* rather than *knowledge-based*, the IEE stops every ransomware strain — past, present, and future — without requiring any knowledge of what the threat is. Antivirus software is no longer needed, and its removal reduces attack surface further. Global ransomware damages reached an estimated **$57 billion annually in 2025**, a figure projected to grow to $275 billion by 2031 absent structural intervention. The IEE represents a viable, low-cost architectural path to ending this trajectory permanently.

Table of Contents

---

1. Introduction: The Broken Default

1.1 The Default OS Security Model Is Inverted

Every mainstream operating system — Windows, Linux, macOS — ships with the same foundational security posture: any code that can be placed on the filesystem can be executed by a sufficiently privileged user. This is the *default allow* execution model, and it is the root cause of the entire modern malware ecosystem.

The implication is not merely theoretical. In practice, the default allow model means that an attacker's only obstacle is getting malicious code onto the target filesystem. Delivery mechanisms — phishing emails, malicious downloads, supply-chain compromises, drive-by browser exploits — are the attacker's entire engineering challenge. Once code reaches disk, the OS cooperates in its execution. The system treats unauthorized code and authorized code identically at the execution layer.

Every major security product category exists as a direct consequence of this single design choice:

- **Antivirus** exists because the OS allows arbitrary code to execute, necessitating a secondary system to evaluate whether that code is malicious.
- **EDR (Endpoint Detection and Response)** exists because code executes before it can reliably be evaluated, requiring post-execution behavioral monitoring.
- **Behavioral heuristics** exist because signature databases are always incomplete relative to the current threat landscape.
- **Sandboxing** exists because execution cannot be prevented, so it must at least be contained.
- **Threat intelligence feeds** exist because the OS grants every unknown executable the same trust as every known one.

This is an inverted security posture. The correct default — deny all unauthorized execution — has been technically feasible for decades. The IEE architecture described in this paper implements it with minimal OS changes.

1.2 The Scale of Failure

The failure of the reactive security model is not a matter of opinion — it is documented in annual statistics that have worsened consistently for twenty-five years. The data for 2025 and 2026 is unambiguous:

Metric	Measured Value	Source
Global ransomware damage cost (2025)	$57 billion annually ($156 million per day)	Cybersecurity Ventures, 2025
Projected global ransomware damage (2031)	$275 billion annually	Cybersecurity Ventures, 2025
Average total breach cost per ransomware incident	$5.08 million	IBM Cost of a Data Breach Report, 2025
Unique victim organizations on leak sites (2025)	7,500+ (58% increase year-over-year)	GuidePoint Security, 2025
Ransomware attacks occurring daily worldwide	Over 2,200	Industry aggregates, 2025

Metric	Measured Value	Source
Ransomware's share of all cybersecurity breaches	37%	Verizon DBIR, 2025
Year-over-year increase in ransomware incidents (2025)	13% (incidents doubled over five years)	Verizon DBIR, 2025
Active ransomware groups tracked (2025)	124	GuidePoint Security, 2025
Median time from intrusion to ransomware deployment	Approximately 5 days	Sophos State of Ransomware, 2025

Table 1.1 — Ransomware and malware impact statistics, 2025. The upward trend across all metrics persists despite two decades of antivirus and EDR investment.

Critically, zero-day ransomware — malware not previously observed in the wild — bypasses every signature-based defense by definition. At the moment of first deployment, no antivirus database contains its signature, no behavioral model has learned to flag it, and no threat intelligence feed has catalogued it. The window of maximum damage is precisely the window during which current defenses are blind.

Despite unprecedented investment in endpoint security products, ransomware incidents continue to grow year-over-year. This is not a failure of individual products — it is the structural consequence of defending a broken default with reactive tooling. More sophisticated reactive tools do not address the root cause.

1.3 The Thesis

The correct fix is not a better antivirus. The correct fix is to change the OS default from *allow all execution* to *allow only authorized execution* — enforced structurally at the kernel level, with immutability guaranteed even against root-level compromise.

This paper describes exactly how to achieve that with three small, targeted kernel modifications and a 25-line plain-text whitelist. No architectural restructuring of the OS is required. The changes are backward-compatible with all authorized software. The resulting security guarantee — no unauthorized code execution, ever — is categorically stronger than any signature-based or behavioral defense can provide.

> ● ***Foundational Principle***
> *"Every reactive security product exists to compensate for the OS permitting what it should have denied. The IEE does not improve the compensatory layer — it eliminates the need for it."*

---

2. The Immutable Execution Environment Architecture

2.1 Core Principles and Invariants

The IEE is defined by four invariants. These invariants are enforced by the kernel — not by userspace security software — and therefore cannot be bypassed by any process, regardless of privilege level, during normal-mode operation.

Invariant 1 — Execution Whitelist

No binary or executable may run unless its absolute path falls within a whitelisted directory. This applies to every user account, including root, SYSTEM, and Administrator. The check occurs inside the kernel, before the process image is mapped into memory. There is no exception for any user, group, or capability flag.

Invariant 2 — Script Path Enforcement

No interpreter — bash, Python, PowerShell, Node.js, Perl, Ruby, PHP, Lua, AWK, or any other registered interpreter — may load or execute a script file unless that script's absolute path falls within a whitelisted directory. This check occurs at the kernel VFS layer and applies to all interpreters identically. It is not possible for a modified or compromised interpreter to bypass a kernel-level access denial.

Invariant 3 — Immutable Whitelist Paths

Every directory listed in the whitelist is read-only during normal operation. No process — regardless of privilege level — may write to, rename, delete, truncate, or modify any file within a whitelisted path. This includes OS binaries, shared libraries,

scripts, and any data files the administrator designates for protection (financial records, configuration stores, backup indices, etc.).

Invariant 4 — Safe Mode Administration Only

The whitelist and all whitelisted content may be modified only when the system is booted into Safe Mode, whose network accessibility is defined by a per-organization configurable policy, sealed in the system's Trusted Platform Module (TPM), and applied at kernel initialization — before any userspace process runs.

2.2 What the Whitelist Contains

The whitelist contains *directory paths*, not individual file hashes. This is the key architectural decision that makes IEE maintainable at production scale. The entire whitelist for a fully operational production system fits comfortably in 15–25 lines of plain text. A representative production whitelist for a Linux system is shown below:

```
# Immutable Execution Environment — Path Whitelist v1.0 # Format: one directory path per line # All subdirectories are included recursively # Lines beginning with # are comments  # ── System executables ──────────────────────────────── /bin /sbin /usr/bin /usr/sbin /usr/local/bin # ── System libraries (loaded by authorized executables) /lib /lib64 /usr/lib /usr/lib64  # ── Authorized scripts only ──────────────────────────── /etc/scripts  # ── System configuration (read-only, non-executable) ─ /etc  # ── Protected data files (unencryptable by definition) /protected/data  # ── Windows equivalent paths (for reference) ─────────── # C:\Windows # C:\Program Files # C:\Program Files (x86) # D:\ProtectedDocuments
```

Adding a new application requires adding exactly one line — its installation directory. OS security patches require *no whitelist changes whatsoever* — paths remain identical across patch cycles, only file content changes, and that content is applied during Safe Mode where writes are permitted. Adding a new protected document or data directory requires one additional line.

This stands in stark contrast to hash-based application allowlisting, which has historically been the primary alternative approach. The comparison is instructive:

Metric	Hash-Based Allowlisting	IEE Path-Based Whitelist
List size	100,000+ individual file hashes	15–25 directory paths
OS security patch	Rehash every updated binary (hours)	No change needed (seconds)
New application install	Hash every new file	Add one directory line

Metric	Hash-Based Allowlisting	IEE Path-Based Whitelist
	individually	
Application update	Rehash all changed files	No change needed
New protected document	Not supported by design	Add one directory line
Admin effort per patch cycle	Hours of hash management	Seconds to zero
False positive risk	High — any missed hash blocks execution	Low — directory-level granularity
Scalability across fleet	Poor — per-system hash variation	Excellent — identical whitelist across fleet

Table 2.1 — IEE path-based whitelist versus hash-based allowlisting. The path-based model achieves equivalent security with orders-of-magnitude less administrative overhead.

---

3. The Three Kernel Modifications

The entire IEE is implemented through three small, targeted modifications to existing kernel subsystems. Each modification inserts a whitelist check at a specific enforcement point in the kernel's existing code path. No architectural restructuring is required. Existing kernel functionality is entirely preserved for all non-whitelisted paths — standard application execution, file creation, script development, and user data access all function normally outside the protected directory set.

Total new code: approximately 80 lines across three files. The whitelist loader module adds approximately 150 additional lines of infrastructure code for radix-tree construction and TPM-sealed configuration parsing.

3.1 Modification 1 — The Execution Gate (execve hook)

The first and most foundational modification is a whitelist check inserted into the kernel's program execution path. In Linux, this is the `do_execve()` function located in `fs/exec.c`. In Windows NT, the equivalent enforcement point is the NT image loader in `ntoskrnl.exe`, accessible via the `PsCallImageNotifyRoutines` callback mechanism or as a Windows Defender Application Control (WDAC) policy extension.

Every time any process attempts to execute a program — regardless of the requesting user's identity, privilege level, or group membership — the binary's resolved absolute path is checked against the whitelist *before* the binary is mapped into memory:

```
/* Linux fs/exec.c — Execution Gate (simplified illustration)  * IEE
Modification 1 of 3  * Approximately 20 lines of new code  * Enforcement
point: before ELF image is mapped into memory  */  int do_execve(const char
*filename,               const char __user *const __user *argv,
const char __user *const __user *envp) {     char *resolved_path;      /*
═══════════════════════════════════════════════════════      * IEE EXECUTION GATE
— inserted before loader entry      *
═══════════════════════════════════════════════════════ */     resolved_path =
iee_resolve_absolute_path(filename);      if (!
iee_whitelist_path_allowed(resolved_path))
{         iee_audit_log(AUDIT_IEE_EXEC_BLOCKED,
"Blocked execution attempt: %s  uid=%d  pid=%d",
resolved_path, current_uid(), current->pid);         kfree(resolved_path);
return -EACCES;   /* Permission denied — process never starts */     }
kfree(resolved_path);     /*
═══════════════════════════════════════════════════════      * End IEE addition —
original loader continues below      *
═══════════════════════════════════════════════════════ */      return
do_execveat_common(AT_FDCWD, filename,                                argv,
envp, 0); }
```

The `iee_whitelist_path_allowed()` function performs a radix-tree prefix match against the loaded whitelist table, constructed at boot time from the TPM-sealed whitelist configuration. The lookup completes in microseconds — a negligible performance penalty on any modern processor. Symlinks are resolved to their canonical absolute path before the check, preventing symlink-based traversal attacks.

If the binary's resolved absolute path does not begin with any whitelisted directory prefix, the syscall returns `-EACCES` immediately. The process never starts. The binary image is never mapped. No exception exists for any user, group, capability (including `CAP_SYS_ADMIN`), or POSIX privilege level. UID 0 (root) receives the identical denial as an unprivileged user.

> ● ***Key Insight — Execution Gate***
>
> *The execution gate fires before the binary is mapped into memory. There is no mechanism by which userspace code — regardless of privilege — can intercept it, work around it, or escalate past it. The gate lives entirely within the kernel, at a layer that precedes all userspace execution. A malicious binary that reaches disk is*

> *inert: it cannot execute, it cannot be loaded, and it cannot be used as a staging vehicle for further compromise.*

3.2 Modification 2 — The Interpreter Script Gate (VFS open hook)

The second modification closes the script execution gap, which is the most significant potential bypass vector in a whitelist-only scheme. Script interpreters — bash, Python, PowerShell, Node.js, Perl, Ruby, PHP, Lua, AWK, and others — are themselves authorized binaries residing in whitelisted paths and therefore permitted to run by the Execution Gate. However, they accept arbitrary external script files as runtime input. Without a second gate, an attacker could deliver a malicious script to any writable, non-whitelisted directory (such as `/tmp`, `~/Downloads`, or an email attachment staging area) and execute it by invoking a trusted interpreter against it.

Rather than patching each interpreter individually — a fragile approach requiring coordination across upstream projects and vulnerable to interpreter replacement or compilation from source — the check is implemented once in the kernel VFS layer, at the `open()` system call:

```
Script Gate Logic — VFS open() hook (pseudocode) IEE Modification 2 of 3
Approximately 30 lines of new code Enforcement point: kernel VFS layer,
before file descriptor is returned  When any process calls: open(path,
O_RDONLY)    1. Is the calling process a registered interpreter?
(Checked against a fixed interpreter process table provisioned at setup)
→ Registered interpreters: bash, sh, dash, python3, python2, node,
perl, ruby, php, lua, awk, tclsh, powershell, pwsh, Rscript    2. Does the
target file have a registered script extension?
→ .sh .bash .py .js .ts .pl .rb .php .lua .awk .ps1 .psm1        .psd1 .tcl .
r .R — configurable at provisioning time    3. Is path within the IEE
whitelist?    Decision matrix:
┌──────────────────────────┬──────────────────────────────────────────┐   │
Interpreter?  Script ext?   | Path in whitelist?  →  Result        |
├──────────────────────────┼──────────────────────────────────────────┤   │
YES           YES           | YES                 →  open() succeeds|   | YES
YES           | NO                  →  -EACCES        |   | YES           NO
| (any)               →  open() succeeds|   | NO            (any)         |
(any)               →  open() succeeds|
└──────────────────────────┴──────────────────────────────────────────┘
Result when blocked: EACCES returned to interpreter.   Interpreter never
reads a single byte of the script.   The script file remains unexecuted
regardless of its content.
```

A malicious `.py`, `.ps1`, `.sh`, or `.js` file dropped into any non-whitelisted directory — `/tmp`, `~/Downloads`, a mounted USB volume, an email attachment directory — cannot be executed by any interpreter on the system, regardless of the invoking user's permissions or privilege level.

> ● ***Key Insight — Script Gate***
> *The script gate is enforced at the kernel VFS layer — below every interpreter. A modified, backdoored, recompiled, or exploit-compromised interpreter running in userspace cannot bypass a kernel-level EACCES return. The gate cannot be defeated by changing the interpreter; it can only be defeated by changing the kernel — which requires Safe Mode access and physical presence.*

3.3 Modification 3 — The Immutability Gate (VFS write hook)

The third modification enforces read-only protection on all whitelisted paths. It is inserted into the VFS write-path dispatcher, through which every file write operation passes regardless of the underlying filesystem type (ext4, XFS, NTFS, APFS, tmpfs, or any other):

```
/* Linux fs/read_write.c — Immutability Gate (simplified illustration)  * IEE
Modification 3 of 3  * Approximately 30 lines of new code  * Enforcement
point: VFS write dispatcher — covers ALL filesystem types  */  ssize_t
vfs_write(struct file *file, const char __user *buf,                  size_t
count, loff_t *pos) {     /*
═══════════════════════════════════════════════════      * IEE IMMUTABILITY
GATE      * ═══════════════════════════════════════════════════ */     if
(iee_path_is_protected(&file->f_path))
{         iee_audit_log(AUDIT_IEE_WRITE_BLOCKED,
"Blocked write to protected path: %s  uid=%d",
dentry_path_raw(file->f_path.dentry, buf, PATH_MAX),
current_uid());         return -EACCES;   /* Denied regardless of caller's
UID */     }     /* ═══════════════════════════════════════════════════
* End IEE addition — original write path continues      *
═══════════════════════════════════════════════════ */      return
__vfs_write(file, buf, count, pos); }  /*  * The same iee_path_is_protected()
check is wrapped around:  *   vfs_rename()    → prevents moving files
within/out of protected dirs  *   vfs_unlink()    → prevents deletion of
protected files  *   vfs_truncate()  → prevents truncation (a write-
equivalent)  *   vfs_chmod()     → prevents permission changes on protected
files  *   vfs_chown()     → prevents ownership changes  *   vfs_setxattr()
→ prevents extended attribute modification  */
```

The caller's identity is completely irrelevant to this check. UID 0 (root on Linux), the SYSTEM account on Windows, and the Administrator account all receive the same `-EACCES` response as an unprivileged standard user. The immutability is enforced by the kernel gate, which executes at a layer no userspace process — regardless of privilege — can reach during normal-mode operation.

For protected data files, the same gate applies unconditionally. Any file in a whitelisted directory — whether a system binary, a shared library, a configuration script, or a financial spreadsheet — cannot be written to, renamed, deleted, truncated, or encrypted by any process during normal operation. Ransomware that reaches the system, receives execution from somewhere, and attempts to encrypt protected data files will receive a stream of `EACCES` errors from every encryption write attempt. The protected data files remain intact.

---

4. Antivirus Is No Longer Needed

4.1 Why Antivirus Exists

Antivirus software exists because the OS default permits arbitrary code execution. In that model, every executable that reaches a system — from any source, via any mechanism — is a potential threat that must be evaluated. AV compensates by maintaining continuously updated threat signature databases, running real-time file system scanners with kernel-level hooks, applying behavioral heuristics to code that has already been allowed to launch, and quarantining identified threats after they have already executed some portion of their payload.

This model is inherently reactive. The AV product is always responding to something the OS already permitted. The window between code arrival and AV detection is the attacker's operational window — and for zero-day ransomware, that window is the entire first deployment cycle, which may encompass thousands of victims before a signature is produced.

4.2 Why Antivirus Becomes Redundant Under IEE

Under the IEE, unauthorized code never executes at all. Each antivirus function maps directly to a structural IEE equivalent that is superior in every measurable dimension:

AV Function	Purpose Under Default-Allow OS	IEE Structural Equivalent
Signature scanning	Match running or stored code against database of known malware hashes	Unauthorized code never executes — there is nothing to match signatures against
Behavioral detection	Monitor what running code does and flag suspicious patterns	Unauthorized code never runs — no behavior to monitor, ever
Heuristic analysis	Flag code with structurally suspicious characteristics (packers, shellcode patterns)	Unauthorized code can never run — heuristics have no input to process
Real-time file scanning	Inspect files as they are opened, written, or downloaded	VFS immutability gate blocks writes to protected paths structurally; non-protected files are inert because they cannot execute
Quarantine	Isolate identified threats to prevent further execution	Threats sit inert in non-whitelisted directories — permanently unable to execute — without requiring identification or quarantine
Zero-day protection heuristics	Attempt to detect novel malware before signatures exist	Zero-day malware is blocked by path check — its novelty is irrelevant; the IEE does not need to know what the threat is
Ransomware rollback	Snapshot and restore encrypted files if ransomware is detected	Protected files cannot be encrypted — rollback capability is unnecessary for whitelisted data

Table 4.1 — Mapping of antivirus functions to IEE structural equivalents. Each AV capability is replaced by a kernel-level structural control that is both stronger and requires no ongoing maintenance.

4.3 AV Removal Reduces Attack Surface

Antivirus products are among the most privileged software on any production system. They run in kernel space, process untrusted and potentially malformed content (compressed archives, packed executables, obfuscated scripts) continuously, parse complex and attacker-controlled file formats, and expose large kernel-mode code surfaces to the rest of the OS. This combination has historically made AV software a prime exploitation target. Multiple major AV vendors have published critical kernel-level CVEs allowing local privilege escalation or remote code execution via the AV scanning engine itself.

Under IEE, removing antivirus software is not a security reduction — it is a net security improvement. The AV kernel driver is itself an attack surface component, and eliminating it removes a class of potential compromise vectors while simultaneously eliminating the need it previously filled.

> ● ***Structural Insight — AV Obsolescence***
> *"The antivirus industry exists because operating systems are broken by default. Fix the OS default — and antivirus becomes as unnecessary as a lock on a vault that has no door. The vault's structural integrity, not the lock, is the correct defense."*

---

5. Fileless and Memory-Based Attacks Are Practically Intractable

5.1 The Theoretical Concern

Fileless attacks — exploiting memory-corruption vulnerabilities in authorized, already-running processes to execute injected shellcode without writing a new binary to disk — represent the one theoretical class of threat that does not require the execution of a new binary through the IEE execution gate. In theory, a sufficiently sophisticated attacker could bypass the execution gate entirely by

injecting code into an already-authorized running process through memory corruption.

This concern is real and deserves rigorous treatment. The following analysis demonstrates that while the theoretical concern is valid, the practical viability of fileless attacks at scale — particularly for ransomware, which requires reliable execution across thousands of heterogeneous, unpredictable target systems — is near zero against a modern hardened system, and is further constrained by IEE even when exploits partially succeed.

5.2 Why This Is Practically Non-Viable at Scale

Successfully exploiting a modern hardened system via memory corruption requires an attacker to chain through multiple independent defensive layers, each of which demands a separate, distinct vulnerability to defeat:

```
Step 1: Discover a memory corruption bug in an authorized process         →
Modern compilers apply: stack protectors (-fstack-protector-strong),
SafeStack, AddressSanitizer (in test builds), bounds checking         →
Memory-safe languages (Rust, Go) eliminate entire bug classes           at
the language level — UAF, buffer overflow, integer overflow         → Kernel
components increasingly rewritten in Rust (Linux 6.x+)  Step 2: Defeat ASLR
(Address Space Layout Randomization)         → Requires a separate
information-leak vulnerability to determine           the randomized base
addresses of loaded modules         → 64-bit ASLR entropy: 2²⁸ to 2⁴⁰
possible layouts depending on arch         → Brute force is infeasible at
these entropy levels         → Leak bug must be found and exploited
independently  Step 3: Defeat stack canaries         → Random value placed
between local variables and return address         → Any stack overflow
corrupts the canary → detected on function return            → process killed
before control reaches attacker code         → Canary value must be leaked
first — another independent vulnerability  Step 4: Defeat DEP / NX (No-
Execute bit on data pages)         → Hardware-enforced: pages marked as data
cannot be executed by the CPU         → Direct shellcode injection into data
memory fails at the hardware level         → Requires Return-Oriented
Programming (ROP) instead:            chaining existing code fragments
(gadgets) from authorized binaries  Step 5: Defeat CFI (Control Flow
Integrity)         → Compiler-inserted checks validate every indirect branch
against           a set of allowed targets computed at compile time         →
Reduces usable ROP gadget set to near zero — most gadgets           are
unreachable because they are not valid indirect call targets  Step 6: Defeat
Intel CET / ARM BTI (hardware shadow stack)         → CPU maintains a
separate hardware shadow stack of return addresses         → Differs from
software shadow stack — not accessible to userspace         → ROP chain
manipulation causes shadow stack mismatch on RET instruction           →
hardware fault → process killed         → ARM Branch Target Identification:
landing pad instructions required           at all valid branch targets —
arbitrary gadget chaining fails at hardware  Step 7: Escape the process
sandbox         → Modern browsers, email clients, document readers run in
low-integrity           sandboxed processes with restricted syscall surfaces
(seccomp-bpf)         → Even a fully successful in-process exploit runs at
```

```
sandbox privilege         → Escaping the sandbox requires a separate kernel
privilege-escalation bug         → The escaped process STILL faces the IEE
immutability gate (see §5.3)
```

Each step listed above requires a separate, independently discovered vulnerability. Nation-state actors with significant, sustained research investment can occasionally chain three or four of these steps for a specific, pre-selected high-value target — after months of reconnaissance and development. Ransomware operators operate on a fundamentally different economic model: they require attacks to work reliably at scale across thousands of heterogeneous, unpredictable, randomly-encountered systems. The exploit chain development cost for reliable cross-system memory exploitation — far exceeding $1 million per chain — makes this economically non-viable for commodity ransomware operations.

5.3 IEE Constrains Even Successful Memory Exploits

Even if an attacker successfully completes every step above — a feat requiring extraordinary resources — the IEE structural gates remain active for all subsequent operations:

- **Cannot write to any whitelisted path** — the immutability gate denies every write attempt at the VFS layer, regardless of the privilege level of the exploited process.
- **Cannot execute a new process outside the whitelist** — any attempt to spawn a persistent executable or download a secondary payload triggers the execution gate.
- **Cannot persist across process death** — injected in-memory shellcode exists only in the address space of the exploited process. When that process exits or crashes, the exploit disappears entirely.
- **Cannot encrypt protected data files** — every encryption write to a whitelisted data directory returns `EACCES` from the immutability gate. Protected files remain intact.
- **Blast radius is bounded** — the attacker's impact is limited to the user-writable, non-whitelisted address space of the currently running process, and vanishes when that process closes.

> ● ***Key Insight — Fileless Containment***
>
> *The IEE does not claim to make memory exploitation impossible — that is a separate research problem addressed by memory safety and hardware mitigations. The IEE claims that even a perfect, complete memory exploit — one that defeats every listed mitigation — still cannot encrypt protected data, persist to disk in a protected path, or spawn new persistent unauthorized processes. The attack dies when the browser tab closes.*

---

6. Safe Mode — Configurable Administration Policy

6.1 The Principle

System administration — updating the whitelist, applying OS security patches, installing new authorized software — must be possible, but must be structurally impossible for remote attackers to reach during normal-mode operation. The IEE achieves this through a Safe Mode whose network accessibility is a per-organization configurable policy, not a fixed implementation detail. Different organizations have fundamentally different threat models and operational constraints; the IEE Safe Mode framework accommodates the full spectrum.

6.2 TCP Stack Removal — The Cleanest Isolation

The most architecturally elegant form of Safe Mode network isolation does not use a firewall. It simply does not load the TCP/IP kernel stack. A firewall is a filtering layer over an active network stack — it can be misconfigured, it contains exploitable code, and the network stack beneath it remains active and reachable through implementation bugs. TCP stack removal takes a different approach: if the networking kernel module is not loaded, there is no code to exploit, no configuration to misconfigure, and no packets that can be processed.

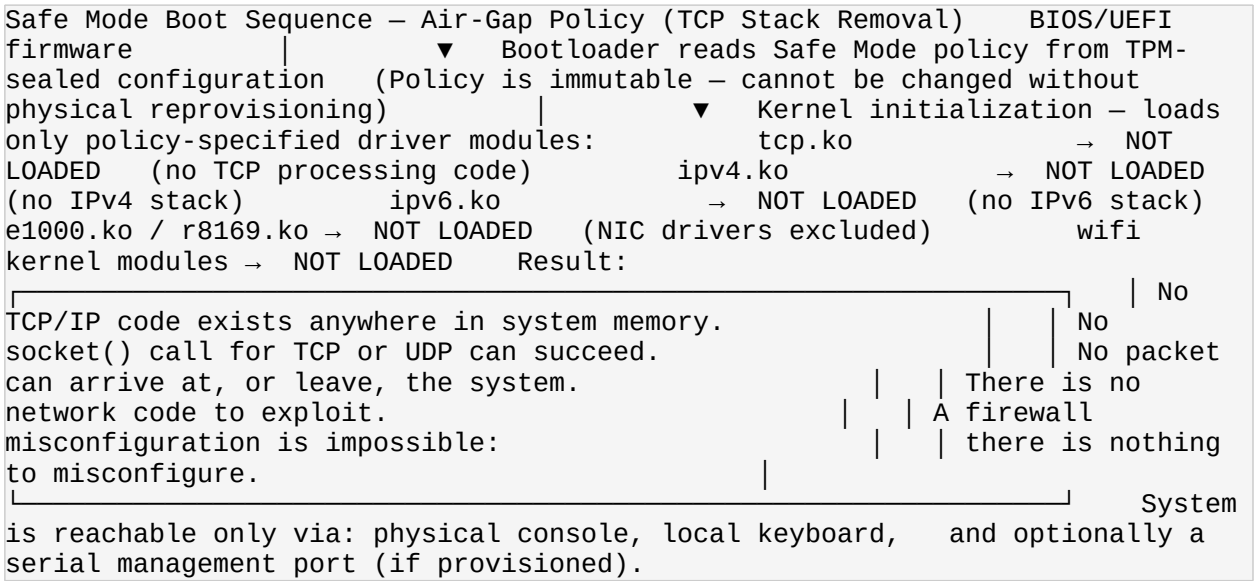

Approach	TCP Stack Loaded?	Network Code in Memory	Misconfiguration Risk	Residual Network Attack Surface
Firewall rules in Safe Mode	Yes — stack runs; rules filter traffic	Full TCP/IP stack present	High — one wrong rule creates exposure	Stack CVEs remain exploitable
IEE TCP stack removal	No — stack module not loaded	Zero bytes of network code	None — nothing to configure	Zero network attack surface

Table 6.1 — Firewall-based isolation versus TCP stack removal. Stack removal eliminates the residual attack surface that firewall-based approaches inherently retain.

6.3 The Safe Mode Policy Spectrum

Not every organization requires or benefits from air-gap Safe Mode. The IEE Safe Mode network policy is configurable at provisioning time, sealed in the TPM, and enforced at kernel initialization before any userspace process runs. The following policy spectrum covers the operational range from classified military infrastructure to cloud-native enterprise environments:

Policy Level	Network Access in Safe Mode	Target Organizations	Enforcement Mechanism
Air-Gap	None — no NIC or TCP drivers loaded	Military, nuclear, classified	Driver exclusion list applied at

Policy Level	Network Access in Safe Mode	Target Organizations	Enforcement Mechanism
		infrastructure	kernel init
LAN-Only	Internal RFC 1918 ranges only; no internet routing	Hospitals, banks, industrial control systems	nftables applied before NIC brought online
IP Whitelist	Only explicitly listed administrator machine IPs	Enterprises, managed service providers	nftables ACCEPT for listed IPs, DROP all others
Subnet + VLAN	Administration VLAN only	Segmented large enterprises	Network-level enforcement at switch + kernel rules
VPN + IP Whitelist	Corporate VPN connection AND whitelisted source IP required	Remote-first, cloud-native organizations	VPN tunnel establishment + IP source verification
Hardware Token + IP	Physical YubiKey press AND IP whitelist match required simultaneously	Financial institutions, defense contractors	Two-factor: physical hardware presence + network source

Table 6.2 — IEE Safe Mode network policy spectrum. Each policy level is sealed in the TPM at provisioning time and cannot be changed without physical access for reprovisioning.

6.4 How the Policy Is Protected

The Safe Mode network policy is stored inside the TPM-sealed whitelist configuration, making it subject to the same immutability as the whitelist itself. A three-tier administrative hierarchy emerges naturally from this design:

```
Administrative Hierarchy — Three Tiers
╔═══════════════════════════════════════════════════════════════════════════╗ ╟ TIER 1 — PROVISIONING                                                    ╟ ╟ When:  One time only, during system commissioning                      ╟ ╟ How:   Physical access to machine required                             ╟ ╟ Sets:  Whitelist paths, Safe Mode network policy,                      ╟ ╟        authorized administrator IPs, TPM binding                       ╟ ╟ Seals: All configuration into TPM. Tier 1 closes permanently.  ╟
╠═══════════════════════════════════════════════════════════════════════════╣ ╟ TIER 2 — SAFE MODE ADMINISTRATION                                         ╟ ╟ When:  Ongoing — for patching, software updates, whitelist edits╟ ╟ How:   Reboot to Safe Mode; authenticate per network policy     ╟ ╟ Can:   Modify whitelist paths, update binaries, install software╟ ╟ Cannot: Change Safe Mode network policy (requires Tier 1 repro) ╟
```

```
╠═══════════════════════════════════════════════════════════════════════════╣ ╠ TIER 3 — NORMAL MODE OPERATION                                          ╠ ╠ When:  All runtime — applications, users, services                      ╠ ╠ How:   Standard login; no special authentication                        ╠ ╠ Can:   Read whitelisted files, write to non-whitelisted paths           ╠ ╠ Cannot: Modify anything in any whitelisted directory                    ╠ ╠         All three kernel gates active and enforcing                      ╠
╚═══════════════════════════════════════════════════════════════════════════╝
```

No remote attacker — regardless of the depth of compromise achieved during normal-mode operation — can reach Tier 1 or Tier 2 without combining physical access to the hardware with knowledge of the Safe Mode credentials. A fully compromised normal-mode session, including root-level compromise, provides zero capability to modify the whitelist, alter the Safe Mode policy, or persist code to protected paths. The attacker's foothold evaporates on the next reboot.

---

7. Comparison with Existing Security Approaches

The IEE architecture does not exist in a vacuum. Multiple existing approaches attempt to address the problem of unauthorized code execution with varying degrees of effectiveness and operational overhead. The table below provides an honest comparative assessment across security-relevant dimensions:

Approach	Mechanism Type	Stops Zero-Days?	Requires Threat Knowledge?	Effective vs. Novel Ransomware?	Admin Overhead	Residual Attack Surface
Antivirus (signature-based)	Reactive — block known-bad by signature match	No — signature must exist first	Yes — signature DB must precede the threat	No — new strains evade until signature produced	High — continuous signature updates, scanning overhead	High — AV kernel driver is itself a significant attack surface
EDR / Behavioral Detection	Reactive — flag statistically suspiciou	Partial — novel behavior patterns may	Yes — must learn baseline "normal"	Partial — novel behavior may not trigger	High — constant tuning, false positive	Medium — privileged agent runs with

Approach	Mechanism Type	Stops Zero-Days?	Requires Threat Knowledge?	Effective vs. Novel Ransomware?	Admin Overhead	Residual Attack Surface
	s runtime behavior	trigger generic rules	to detect "suspicious"	tuned rules; false negatives common	management, agent updates	high system access
Application Allowlisting (hash-based)	Proactive — block execution of any binary not in approved hash list	Yes — unknown hashes are blocked regardless of novelty	No — no threat knowledge required	Yes — ransomware binary has no approved hash	Very high — every patch cycle requires rehashing all updated binaries	Low — but hash management overhead causes frequent abandonment
WDAC + HVCI (Windows)	Proactive — hypervisor-enforced code signing policy	Yes — unsigned code is blocked	No — code signing check is structural	Yes — unsigned ransomware binaries are blocked	Medium — certificate management; vendor signing requirements	Low — hypervisor layer provides strong isolation; data files unprotected
IEE (this paper)	Structural — OS execution default changed at kernel level	Yes — path check is independent of threat content or signature	No — whitelist requires no knowledge of threats	Yes — all ransomware variants, past and future, are blocked structurally	Near-zero — 25-line text file; no per-patch changes required	Near-zero — AV removed, data files protected, minimal kernel additions

Table 7.1 — Comparative analysis of security approaches across key dimensions. IEE achieves the strongest security guarantee with the lowest operational overhead of any mechanism assessed.

---

8. Implementation Guide

8.1 Linux Implementation

The IEE kernel modifications for Linux target well-defined, stable internal functions with known calling conventions and minimal upstream churn. The recommended implementation path is as a Linux Security Module (LSM), which provides a clean hook framework, enables upstream submission without touching core kernel files, and allows enable/disable at compile time via `CONFIG_SECURITY_IEE`.

Modification	Kernel File	Target Function	Lines Added	LSM Hook
Execution Gate	`fs/exec.c`	`do_execve()`	~20	`security_bprm_check()`
Script Gate	`fs/namei.c`	`vfs_open()`	~30	`security_file_open()`
Immutability Gate (write)	`fs/read_write.c`	`vfs_write()`	~15	`security_file_permission()`
Immutability Gate (metadata)	`fs/attr.c`, `fs/namei.c`	`vfs_rename()`, `vfs_unlink()`, `vfs_truncate()`	~15	`security_inode_rename()`, `security_inode_unlink()`
Whitelist Loader Module	`security/iee/iee_whitelist.c` (new)	Boot-time module init	~150	LSM init hook + radix tree construction
Safe Mode Boot Parameter	Kernel command line + `init/main.c`	Early init detection	~10	Kernel boot parameter parsing

Table 8.1 — Linux kernel modification map. All modifications are expressible as LSM hooks, enabling clean upstream submission without modification of core VFS functions.

8.2 Windows Implementation

On Windows, the IEE maps to well-documented kernel extension mechanisms that do not require modification of signed OS binaries:

- **Execution Gate:** Implemented via `PsSetLoadImageNotifyRoutine` callback in a kernel-mode driver, or extended as a Windows Defender Application Control (WDAC) supplemental policy with custom path rules. HVCI-compatible implementation requires the driver to be WHQL-signed.
- **Script Gate:** Kernel callback on `NtCreateFile` / `NtOpenFile` via a minifilter driver pre-operation callback, filtered to registered interpreter process IDs and script file extensions. Minifilter altitude: 320000–329999 (FSFilter Anti-Virus range provides appropriate priority).
- **Immutability Gate:** File system minifilter driver registered for `IRP_MJ_WRITE`, `IRP_MJ_SET_INFORMATION` (covers rename, delete, truncate), and `IRP_MJ_CREATE` with write-access flags. Path comparison against whitelist loaded at driver initialization.
- **Whitelist Storage:** Stored in a UEFI authenticated variable, TPM-sealed and bound to PCR values representing the boot chain. Loaded by the minifilter driver at driver initialization, before any user process starts.
- **Safe Mode:** Windows Safe Mode boot entry created via `bcdedit`; custom service list excludes TCP/IP stack services (`Tcpip`, `Ndis`, NIC driver services) for air-gap policy. Service exclusion list is sealed in UEFI variable alongside whitelist.

8.3 Deployment Checklist

#	Deployment Step	Verification Test	Status
1	Define whitelist directory paths for all authorized software	Review whitelist for completeness against installed software inventory	Provisioning
2	Add critical data directories to whitelist as protected (read-only) paths	Confirm data directories are listed in whitelist	Provisioning

#	Deployment Step	Verification Test	Status
3	Register all script interpreters in the interpreter recognition list	Confirm all interpreter binaries are identified by path	Provisioning
4	Choose Safe Mode network policy appropriate for the organization's threat model	Policy level documented and approved by security team	Provisioning
5	Provision TPM with sealed whitelist and network policy configuration	TPM PCR values confirmed; sealed blob verified	Provisioning
6	Boot to Safe Mode; install and verify all authorized software	All required applications run correctly in Safe Mode	Safe Mode
7	Reboot to Normal Mode; validate all three enforcement gates are active	Kernel log confirms IEE gates loaded and active	Normal Mode
8	Remove antivirus software (no longer required)	AV agent and kernel driver confirmed uninstalled	Normal Mode
9	Test: attempt to execute binary from `~/Downloads/`	Must return `EACCES`; process must not start	Validation
10	Test: attempt to write to `/usr/bin` or `C:\Windows` in Normal Mode	Must return `EACCES`; file system unchanged	Validation
11	Test: attempt to run a script from `/tmp/` via Python or bash	Must return `EACCES` from open(); interpreter prints permission denied	Validation
12	Test: boot to Safe Mode; verify network policy enforces correctly	For air-gap: no network interfaces up, no TCP sockets possible	Validation

Table 8.2 — IEE deployment checklist. Steps 1–5 occur at provisioning time and cannot be repeated without physical access. Steps 6–8 occur in Safe Mode. Steps 9–12 are validation tests in Normal Mode.

---

9. Real-World Proof of Concept — Existing Systems That Validate the Model

The IEE is not a theoretical proposal without precedent. Three large-scale production systems already implement close analogues of its core principles — and their security records validate the architectural approach with data spanning hundreds of millions of deployed devices.

9.1 Chrome OS — Verified Boot and Read-Only Root Filesystem

Chrome OS implements a structural analogue to the IEE immutability gate: the root filesystem is mounted read-only and verified block-by-block via `dm-verity`, a Linux kernel feature that computes cryptographic hashes over each 4KB block and checks them against a signed hash tree on every read. A root-level compromise cannot persistently modify the root filesystem — even with full root privileges — because the underlying block device is write-protected at the firmware layer and the hash tree is signed by Google's key.

System updates are staged to an inactive partition in the background during normal operation — exactly the IEE Safe Mode update model, implemented at partition granularity rather than directory granularity. The system switches to the updated partition on reboot after TPM-verified validation. Chrome OS has never experienced a ransomware infection affecting the system partition across its entire deployment history — a record attributable directly to this architectural choice, not to antivirus software (which Chrome OS does not run).

9.2 iOS — Signed Code Execution and Read-Only System Volume

iOS enforces that only Apple-cryptographically-signed code may execute on the platform. The signing check occurs at the kernel level via the AMFI (Apple Mobile File Integrity) kernel extension, which implements a check conceptually identical to the IEE execution gate. The system volume (`/`) is mounted read-only, cryptographically sealed with a signed hash, and — since iOS 15 — stored on a physically separate NAND partition from user data, making cross-contamination structurally impossible.

No process — including those with the highest available privilege equivalents — can write to the system volume at runtime. iOS has never experienced a mass ransomware outbreak despite deployment across more than one billion active devices worldwide. The absence of mass-market ransomware on iOS is a direct consequence of structural code execution control, not of any reactive security product.

9.3 Windows WDAC + HVCI — Hypervisor-Protected Code Integrity

Windows Defender Application Control (WDAC) combined with Hypervisor-Protected Code Integrity (HVCI) represents the closest existing Windows analogue to the IEE execution gate. WDAC enforces code signing and path-based execution policies. HVCI operates from the hypervisor layer — below the OS kernel itself — ensuring that even a fully compromised Windows kernel cannot disable the code integrity checks. Driver code is verified against a policy before being loaded into kernel address space.

The primary gap between WDAC + HVCI and the full IEE specification is the absence of: (1) the interpreter script gate (WDAC does not enforce script execution paths at the VFS level), (2) data file immutability for user documents, and (3) a configurable Safe Mode network policy sealed in the TPM. The IEE architecture generalizes and extends these proven production approaches, adding the three missing capabilities to form a unified, cross-platform specification built from the same structural principles that already work at production scale.

---

10. Conclusion

10.1 Summary of Capabilities

The Immutable Execution Environment addresses ransomware and malware at their root cause: the operating system default of unrestricted code execution. Through three surgical kernel modifications totaling approximately 80 lines of new code, a 25-line path-based whitelist, and a configurable Safe Mode administration policy, the IEE achieves the following structural guarantees simultaneously:

Guarantee	Mechanism	Coverage
Structural impossibility of unauthorized binary execution	Execution Gate (Kernel Modification 1)	All users, all privilege levels, all future ransomware variants
Script execution restricted identically to binary execution	Script Gate (Kernel Modification 2)	All registered interpreters, all script extensions
Protected data files permanently unencryptable	Immutability Gate (Kernel Modification 3)	All whitelisted data directories, all write operations
Antivirus eliminated as a security dependency	Structural prevention renders AV functions redundant	All AV function categories; AV removal improves security
Fileless memory exploits contained even when successful	Immutability Gate blocks persistence; Execution Gate blocks process spawning	All injected shellcode payloads; attack dies with process
Administration flexibility matching organizational threat models	Configurable Safe Mode network policy sealed in TPM	Air-gap to VPN+IP across six policy levels

Table 10.1 — IEE structural guarantees summary. Each guarantee is enforced by the kernel at a layer no normal-mode userspace process can reach.

10.2 The Fundamental Shift

Current security products — antivirus, EDR, behavioral detection, threat intelligence — all operate *reactively*: they observe what the OS has already allowed to happen

and attempt to identify threats within that execution. This is a fundamentally inverted posture. The defender must be correct every time; the attacker needs to succeed once. Every new malware sample, every novel obfuscation technique, every zero-day deployment resets the defender's advantage to zero.

The IEE operates *proactively*: it changes what the OS allows to happen, so that unauthorized threats never receive execution opportunity at all. The defender's advantage is permanent and does not degrade with the emergence of new threats. A ransomware strain written and deployed in 2035, using techniques not yet imagined, is stopped by the same 25-line whitelist deployed in 2026 — because it is not in the whitelist, and the IEE does not need to know anything about it to stop it.

> *"This is not a better antivirus. It is the replacement of the condition that made antivirus necessary in the first place."*

10.3 Call to Action

OS vendors, kernel security maintainers, security researchers, enterprise architects, and policymakers are called to evaluate the IEE model for incorporation into mainstream operating systems and security frameworks at the following levels:

- **OS Vendors (Microsoft, Linux kernel community, Apple):** Evaluate IEE kernel modifications for mainline inclusion as an optional LSM (Linux), a WDAC extension policy type (Windows), or a system integrity policy extension (macOS). The modifications are small, well-contained, and backward-compatible with all authorized software.
- **Security Researchers:** Develop formal verification of the three invariants, evaluate interaction with eBPF, FUSE, and kernel module loading, and produce a reference implementation against a current LTS kernel.
- **Enterprise Architects and CISOs:** Pilot the IEE model in air-gapped or high-security segments where the whitelist can be defined with precision,

measure operational overhead, and validate the security guarantee in production conditions.

- **Policymakers and Standards Bodies:** Consider IEE-equivalent structural controls as a required baseline for critical infrastructure operators, analogous to the ASD Essential Eight's Application Control recommendation but with the stronger structural guarantees described herein.

The cost of implementation is measured in days of kernel engineering work and weeks of enterprise piloting. The benefit — the permanent structural elimination of ransomware and the majority of the malware threat landscape — is measured in the tens of billions of dollars annually currently absorbed as ransomware damage globally. The asymmetry strongly favors action.

> ***Final Summary***
> *Three kernel modifications. Eighty lines of code. A 25-line whitelist. No antivirus. No EDR. No signatures. No threat databases. Ransomware stopped — structurally, permanently, and for every future variant that will ever be written. The technology is ready. The question is whether the industry is willing to fix the default rather than sell another layer of compensation for it.*

References

- **[1]** Linux kernel source — `fs/exec.c` (`do_execve()`), `fs/read_write.c` (`vfs_write()`), `fs/namei.c` (`vfs_open()`). The Linux Kernel Organization. Available: kernel.org/pub/linux/kernel. Accessed May 2026.
- **[2]** Linux Security Module (LSM) framework documentation. Linux Kernel Documentation Project. Available: kernel.org/doc/html/latest/security/lsm.html. Accessed May 2026.
- **[3]** Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., and Lear, E. "Address Allocation for Private Internets." IETF RFC 1918, February 1996.

- **[4]** Souppaya, M. and Scarfone, K. "Guide to Application Whitelisting." NIST Special Publication 800-167. National Institute of Standards and Technology, October 2015.
- **[5]** Chen, L., Landfermann, R., Löhr, H., Rohe, M., Sadeghi, A.-R., and Stüble, C. "Platform Firmware Resiliency Guidelines." NIST Special Publication 800-193. National Institute of Standards and Technology, May 2018.
- **[6]** Australian Signals Directorate (ASD). "Essential Eight Maturity Model." Australian Cyber Security Centre. Application Control listed as Mitigation Strategy #1. Available: cyber.gov.au/resources-business-and-government/essential-cyber-security/essential-eight. Accessed May 2026.
- **[7]** Cybersecurity and Infrastructure Security Agency (CISA). "Known Exploited Vulnerabilities Catalog." U.S. Department of Homeland Security. Available: cisa.gov/known-exploited-vulnerabilities-catalog. Accessed May 2026.
- **[8]** Microsoft. "Windows Defender Application Control (WDAC) and AppLocker Overview." Microsoft Learn Documentation. Available: learn.microsoft.com/en-us/windows/security/application-security/application-control. Accessed May 2026.
- **[9]** Microsoft. "Hypervisor-Protected Code Integrity (HVCI)." Microsoft Security Documentation. Available: learn.microsoft.com/en-us/windows-hardware/design/device-experiences/oem-hvci-enablement. Accessed May 2026.
- **[10]** Google Chromium Project. "Verified Boot Design Documentation." Chromium OS Design Documents. Available: chromium.googlesource.com/chromiumos/docs/+/HEAD/verified_boot.md. Accessed May 2026.
- **[11]** Apple Inc. "Apple Platform Security Guide." Apple Platform Security. 2026 edition. Available: support.apple.com/guide/security. Accessed May 2026.
- **[12]** Trusted Computing Group. "TPM Library Specification, Family '2.0,' Level 00, Revision 01.59." TCG Published Specification, November 2019. Available: trustedcomputinggroup.org/resource/tpm-library-specification.
- **[13]** Intel Corporation. "Control-flow Enforcement Technology (CET) Specification." Document Number 334525-003US. Intel Architecture Instruction Set Extensions Programming Reference, 2020.
- **[14]** Linux kernel documentation. "dm-verity: Device-Mapper Verity Target." Available: kernel.org/doc/html/latest/admin-guide/device-mapper/verity.html. Accessed May 2026.

- **[15]** Cybersecurity Ventures. "2025 Cybersecurity Almanac: 100 Facts, Figures, Predictions and Statistics." Cybersecurity Ventures, 2025. Cited in: annual global ransomware damage cost $57 billion; projected $275 billion by 2031.
- **[16]** IBM Security. "Cost of a Data Breach Report 2025." IBM Corporation, 2025. Average total cost of ransomware breach: $5.08 million per incident.
- **[17]** Verizon. "2025 Data Breach Investigations Report (DBIR)." Verizon Business, 2025. Ransomware share of breaches: 37%; number of incidents doubled over five years; 64% of victims refused to pay.
- **[18]** GuidePoint Security. "Ransomware Annual Report 2025." GuidePoint Security Research Team, 2025. Unique victim organizations: 7,500+ (58% year-over-year increase); 124 active ransomware groups.
- **[19]** Sophos. "State of Ransomware 2025." Sophos Ltd., 2025. Average recovery cost (excluding ransom): $1.53 million; median time to deployment: approximately 5 days.
- **[20]** Sekar, V., Krishnamurthy, B., et al. "Control-Flow Integrity: Precision, Security, and Performance." ACM SIGSAC Conference on Computer and Communications Security (CCS), 2005. Foundation of CFI defenses referenced in Section 5.
