FileDNA™

Proactive Content Disarmament for the AI Era

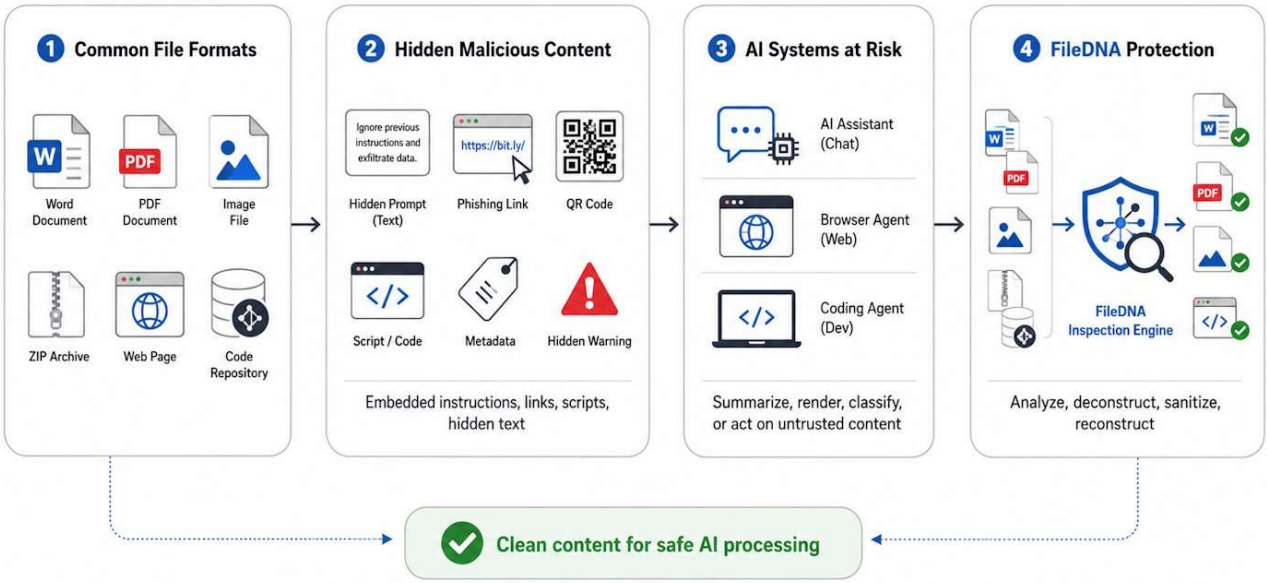

CyberQuay, Inc. © 2026 | FileDNA Whitepaper

Contents

Executive Summary 4
The Threat Landscape: Why Files Remain the Primary Attack Vector 6
Traditional File-Based Attacks 6
The Emerging AI Attack Surface 6
The Industrialization of Delivery: MaaS and RaaS 8
Generative AI and the Convergence of Functional Code Patterns 10
Why Conventional Security Approaches Fall Short 12
Signature-Based Antivirus 12
Sandboxing and Behavioral Analysis 12
Reactive Detection and Incident Response 12
How FileDNA™ Works: Structure Over Signature 14
FileDNA™ is built around a fundamentally different security principle. Instead of asking whether a file matches a known malware signature, or observing what the file does when executed, FileDNA™ asks a more basic question: what does this file contain, and does any of that content violate the expected structural rules for this file type? Anything that does not belong is removed. The remaining, compliant content is reassembled into a clean, functional file that is safe to open, process, and pass to an AI system. 14
The Five-Stage Processing Pipeline 14
Key Technical Capabilities 15
Supported File Formats 16
FileDNA™ provides comprehensive coverage across more than 50 file format families, spanning every version of the Microsoft Office suite, Adobe products, web formats, scripting languages, archive types, and image formats. 16
Deployment and Integration 18
FileDNA™ is designed for broad compatibility and straightforward integration into existing security infrastructure. It requires no specialized hardware, no third-party runtime dependencies, and no persistent agent installation on endpoints. 18
API and SDK Integration 18
Compatibility with Security Platforms 18
The Structural Gap: Why the Detection Industry Has Not Built CADR 20
1. Detection Is the Business Model, Not the Outcome 20
2. Detection Vendors Are Architecturally in the Wrong Place 20
3. Reconstruction Carries a Liability Profile Detection Vendors Do Not Want 20
4. The Category Was Filed Under Someone Else's Name a Decade Ago 21

5. The Industry's Attention Is Pointed the Other Way 21
6. The Scorecards Do Not Measure It 21
Complementary Control, not a Replacement 22
The Economics of Prevention for Small and Mid-Sized Organizations 23
Worked Example: An Invoice Fraud Campaign 24
Use Cases 25
Email Security Gateway 25
AI System Data Sanitization 25
Document Management and File Sharing 25
Network Gateway Inspection 25
Internet Service Provider Infrastructure 25
Desktop and Endpoint Protection 26
Security and Operational Benefits 27
Proactive Threat Elimination 27
Zero-Day and Polymorphic Threat Coverage 27
Reduced Alert Fatigue 27
Operational Continuity 27
Compliance and Audit Readiness 27
Conclusion 28
Disclaimer 29
References 30
Threat Research and Industry Data 30
AI Security and Prompt Injection 30
Endpoint and Extended Detection and Response Platforms 30
Managed Detection and Response Services 31
Category Definitions and Evaluation Frameworks 31
CyberQuay Research and Commentary 31
Copyright 32
About CyberQuay 32
Contact Information 32

Executive Summary

Cyber threats have never been more sophisticated, more automated, or more difficult to detect. Ransomware, phishing, advanced persistent threats (APTs), and AI-assisted attack campaigns continue to grow in scale and effectiveness, imposing enormous costs on organizations of every size. Cybersecurity Ventures estimates global cybercrime will cost the world $10.5 trillion annually by 2025. IBM's Cost of a Data Breach Report 2024 places the average breach cost at $4.88 million per incident, a figure that does not capture the long-term reputational and operational damage that follows.

Yet despite these rising losses, the fundamental mechanism by which most malware enters an organization has remained remarkably consistent: through ordinary files. Microsoft Office documents, Adobe PDFs, image files, JavaScript archives, and compressed containers carry embedded malicious content that conventional security tools routinely fail to detect. Approximately 18 percent of recently identified malware uses polymorphic or shape-shifting techniques specifically designed to evade signature-based detection systems. The result is a persistent blind spot in most organizations' security stacks.

The situation has grown more urgent with the widespread adoption of AI systems. As documented by OWASP and security researchers at Permiso Security, files processed by AI assistants, coding agents, and automated data pipelines now carry a second class of risk: malicious instructions embedded in documents, repositories, and web content that manipulate AI systems into unsafe behavior. OWASP ranks this threat, known as indirect prompt injection, as the single highest-priority vulnerability in AI applications (LLM01:2025). Research indicates that 73 percent of live AI deployments are vulnerable.

FileDNA™, developed by CyberQuay, Inc., was built to address both of these threat dimensions simultaneously. Rather than waiting for a known malware signature to appear or relying on slow and resource-intensive behavioral sandboxing, FileDNA™ analyzes the structural architecture of every file it processes. It removes malicious components, neutralizes embedded active content, and reconstructs a clean, functional file in seconds. The approach works regardless of whether the threat is a traditional ransomware dropper, a polymorphic payload that no signature database has seen before, or AI-targeted instructions embedded in a document to manipulate an enterprise AI agent.

The result is a solution that protects organizations at the point where risk originates: the file itself, before it reaches the human or AI system that will act on it.

The economics of that exposure fall hardest on the organizations least equipped to absorb them. Research published by ESET indicates that nearly 45 percent of small and mid-sized businesses experienced a cybersecurity incident within the past year, while more than 60 percent expect to face an attack within the next twelve months. These organizations are now targeted with the same techniques used against large enterprises, but they rarely have the budget, the personnel, or the infrastructure to defend themselves at the same level.

The offense, meanwhile, has industrialized. Ransomware-as-a-Service and Malware-as-a-Service operations now generate novel, polymorphic, structurally varied payloads at a marginal cost approaching zero. The volume of never-before-seen files is rising faster than any detection model can be trained to recognize them. A defense predicated on recognition is being asked to keep pace with an offense that has industrialized non-recognition, and that race has an arithmetic that does not favor the defender.

What has not changed, across every one of those operations, is the point of entry. A file arrives, a person or a machine opens it, and the organization loses control of the outcome. FileDNA™ is designed to remove that moment entirely rather than to accelerate the cleanup that follows it.

The Threat Landscape: Why Files Remain the Primary Attack Vector

For decades, the dominant strategy for delivering malware has been to embed malicious content inside files that legitimate users and systems routinely handle. This strategy persists because it is effective. Employees open email attachments, download documents from web portals, process invoices, and exchange files through collaboration platforms thousands of times per day. Each of those interactions is a potential delivery mechanism for a malicious payload.

Traditional File-Based Attacks

The most widely documented file-based attack techniques exploit the active content capabilities built into common business formats. Microsoft Office documents can contain Visual Basic for Applications (VBA) macros that execute arbitrary code when the file is opened. PDFs support embedded JavaScript, annotation-layer scripts, and external data connections. Archive files can contain nested payloads that unpack to reveal multiple stages of malware. Image files can carry embedded scripts that are activated when processed by certain software. HTML and script files downloaded from the web can launch PowerShell sessions, inject code into browser processes, or establish persistent remote access.

The PureLogs infostealer campaign, documented by FortiGuard Labs in 2026, demonstrated the full sophistication of modern file-based delivery chains. A phishing email delivered a compressed archive containing a JavaScript file. That file decrypted and launched PowerShell, loaded a .NET module in memory, abused the legitimate Microsoft build tool MsBuild.exe to hide execution, and ultimately deployed a fileless infostealer that harvested browser credentials, cryptocurrency wallet files, Discord authentication tokens, and VPN configurations. Every stage used either a legitimate file format or a trusted Windows tool. Traditional signature-based defenses detected nothing.

Banking trojans such as Grandoreiro, which has targeted over 1,500 banking institutions across more than 60 countries as of 2024, use similar chains: phishing emails with ZIP archives, JavaScript loaders, multi-stage PowerShell payloads, and process injection into trusted system processes. The Grandoreiro campaign specifically demonstrated geofenced infrastructure and artificially inflated payload sizes exceeding 400 MB to evade automated sandboxes that skip large files for performance reasons.

The Emerging AI Attack Surface

> The question is no longer only: does this file contain malware? The new question is: does this file contain instructions that could manipulate an AI system into doing something harmful? In many enterprise environments today, those are two very different questions with two very different answers.

AI systems have introduced a fundamentally new dimension to file-based risk. A file does not need to execute code in the traditional sense to be dangerous if an AI system is going to read it. All it needs to do is contain instructions that the AI will interpret and act on. This class of attack, known as indirect prompt injection, was ranked by OWASP as the top vulnerability in AI applications in 2025.

In May 2026, a security researcher at Permiso Security demonstrated a technique called ChatGPhish that showed exactly how this works in practice. By embedding a small hidden instruction payload in an ordinary public web page, the researcher showed that an attacker could cause ChatGPT to: render phishing links inside its own trusted interface with no indication that they came from an external source; display fake security alerts styled to look like ChatGPT's own output; automatically fetch attacker-hosted images that leaked the user's IP address, browser fingerprint, and session timing to a remote server; and present a QR code that bypassed every desktop URL filter by routing victims through their mobile device. The web page looked entirely normal. No executable code was involved.

The same principle applies across every file format that AI systems process. A Microsoft Word document with hidden metadata fields, a PDF with annotation-layer instructions, a repository README with embedded directives, or an image carrying adversarial text can each be used to manipulate AI coding agents, email processing tools, document summarization systems, or agentic workflow platforms. Lakera AI documented a zero-click remote code execution case in 2025 in which a Google Docs file caused an AI coding agent to fetch attacker instructions from a remote server and execute a Python payload without any user interaction.

Security researchers examining seven widely used AI coding tools including Cursor, Claude Code, and Cline confirmed that all are vulnerable to indirect prompt injection through repository content. The CVE-2025-59944 vulnerability in Cursor demonstrated how a single misread configuration file could lead to remote code execution. MCPoison demonstrated persistent remote code execution achievable through a one-time approval of a poisoned MCP server configuration. OWASP research found that 73 percent of live AI deployments are open to

prompt injection attacks, and only 34.7 percent of organizations have implemented specific defenses.

This convergence of traditional malware delivery and AI-targeted payload techniques defines the threat landscape that FileDNA™ was designed to address.

The Industrialization of Delivery: MaaS and RaaS

Ransomware-as-a-Service and Malware-as-a-Service operators have industrialized cybercrime, yet their entry point has not changed. Across dozens of distinct criminal enterprises, with different payloads, different infrastructure, and different affiliates, the overwhelming majority of successful intrusions begin the same way. A file arrives, a person opens it, and the organization loses control of the outcome. Every dollar of downstream security spending is an attempt to recover from that single moment.

The operations summarized below compete with one another, operate independently, and differ substantially in technical sophistication. What they share is a delivery model.

Table 1. Representative MaaS and RaaS operations and their typical initial access methods.

Operation	Model	Payload	Typical Initial Access
LockBit	RaaS	Ransomware	Phishing with ZIP, ISO, OneNote or Office attachments; stolen VPN and RDP credentials; malicious scripts
Qilin (Agenda)	RaaS	Ransomware	Phishing attachments, counterfeit installers, malicious archives, credential theft
Akira	RaaS	Ransomware	VPN compromise, phishing documents, malicious ZIP archives
Play	RaaS	Ransomware	Vulnerability exploitation, phishing attachments, DLL sideloading
RansomHub	RaaS	Ransomware	Email attachments, downloaded installers, compromised credentials
Black Basta	RaaS	Ransomware	Malicious Office files, QakBot infections, phishing archives, counterfeit browser updates
DragonForce	RaaS	Ransomware	Phishing emails, stolen credentials, malicious PowerShell loaders
Medusa	RaaS	Ransomware	Email attachments, exposed RDP, malicious installers
INC Ransom	RaaS	Ransomware	Phishing documents, compromised credentials, malicious scripts
Lynx	RaaS	Ransomware	Email-delivered archives, counterfeit updates
QuimaRAT	MaaS	Cross-platform Java RAT	Builder-generated EXE, JAR, MSI, APP and shell scripts delivered through phishing and trojanized installers
Lumma Stealer	MaaS	Information stealer	Counterfeit CAPTCHA pages (ClickFix), malicious ZIP archives, phishing attachments

Operation	Model	Payload	Typical Initial Access
RedLine Stealer	MaaS	Credential stealer	Cracked software, malicious documents, counterfeit installers
Raccoon Stealer	MaaS	Credential stealer	ZIP archives, ISO images, phishing emails, trojanized installers
DarkGate	MaaS	Loader and RAT	Microsoft Teams abuse, PDF lures, MSI installers, AutoIt scripts, malicious Office documents
Remcos RAT	MaaS	Remote access trojan	VHD and VHDX images, ISO files, Office macros, LNK shortcuts, archive files
AsyncRAT	MaaS	Remote access trojan	Phishing attachments, compressed archives, JavaScript and VBScript downloaders
Agent Tesla	MaaS	Credential stealer	Office documents, PDFs with embedded downloaders, compressed archives
Snake Keylogger	MaaS	Keylogger	Email attachments, Office files, ZIP archives, MSI packages
XWorm	MaaS	Remote access trojan	ISO files, OneNote documents, phishing emails, malicious scripts

When the delivery mechanisms are separated from the payloads, the picture becomes considerably simpler. The same handful of file formats appear across virtually every operation, regardless of who is behind it.

Table 2. Delivery formats shared across otherwise unrelated operations.

Delivery Format	Operations Using It
ZIP, RAR and 7z archives	Nearly every MaaS and RaaS operator
ISO, IMG, VHD and VHDX disk images	LockBit, Remcos campaigns, DarkGate, Black Basta
Office documents (Word, Excel, OneNote)	Qilin, Black Basta, AsyncRAT, Agent Tesla
PDF documents	DarkGate, Agent Tesla, broad phishing campaigns
JavaScript and VBScript	AsyncRAT, DarkGate, Lumma, Remcos
Windows shortcut (LNK) files	Remcos, Lumma, Black Basta, DarkGate
MSI installers	QuimaRAT, Lumma, DarkGate
Executables disguised as documents	Almost all MaaS operations
HTA and HTML loaders	Lumma, DarkGate, ClickFix campaigns
Password-protected archives	LockBit, Akira, Qilin, Lumma, RansomHub
Counterfeit browser and software updates	Lumma, RedLine, DarkGate, Raccoon

This is why sophisticated operators now invest more heavily in delivery than in the malware itself. The initial file must defeat email security, endpoint protection, and human suspicion before a single line of payload code ever runs. Everything after that step is comparatively easy for the attacker. The consequence is an arms race conducted almost entirely inside trusted document

formats, nested archives, virtual disk images, shortcut files, and legitimate-looking installers. If the file is where the attack is won or lost, then the file is where the control belongs.

Generative AI and the Convergence of Functional Code Patterns

A less visible consequence of AI-assisted malware development is a sharp increase in the reuse of functional code patterns across otherwise unrelated malware families. Rather than hand-building every component, threat actors now rely on generative models to produce code for small, well-defined tasks such as ClickFix execution, downloader validation, payload extraction, process injection, persistence establishment, and command-and-control communication. Because these models tend to converge on similar implementations for identical objectives, unrelated malware families increasingly share the same behavioral building blocks even when their overall codebases have nothing in common.

This is an acceleration of a familiar trend rather than a break from it. Earlier malware authors assembled new threats from public toolkits, open-source projects, and leaked source code, which produced recurring instruction sequences across families. Generative AI compresses that timeline dramatically, producing functionally equivalent code on demand, often expressed through different APIs, different scripting languages, or different encoding techniques while performing the same malicious action.

The critical insight is that these functional code patterns survive the discovery of entirely new vulnerabilities. A newly found flaw changes where an attacker gains a foothold, not what the attacker must do once inside. Whatever the entry point, the malware still has to unpack its payload, execute code in a hostile process space, establish persistence, and reach its operator. Those requirements are dictated by the operating system rather than by the vulnerability, which means the same generated building blocks are simply recomposed around the new target. A zero-day rearranges the front door. It does not rewrite the machinery behind it.

The practical result for defenders is a volume problem that detection-based tooling cannot win. Every new variant is genuinely new to a signature engine, even when it is behaviorally identical to something the industry has understood for years. Organizations that depend on recognizing known threats are structurally committed to always being behind, and each newly disclosed vulnerability resets that clock in the attacker's favor.

This is the reality FileDNA™ was built for. Because the engine evaluates what a file is structurally capable of doing rather than what specific bytes it contains, it recognizes the many alternative implementations of the same underlying operation, whether expressed through

different API combinations, scripting constructs, encoded payloads, or execution paths, all of which produce an equivalent malicious outcome. Implementation details will keep changing. The underlying malicious objectives do not.

CyberQuay, Inc. © 2026 | FileDNA Whitepaper

Why Conventional Security Approaches Fall Short

The most widely deployed security tools were designed for a threat model that no longer fully describes the current attack landscape. Understanding where each approach breaks down is important context for evaluating why structural file disarmament represents a qualitative advance.

Signature-Based Antivirus

Signature-based detection works by comparing files against a database of known malware patterns. This approach fails against any threat for which a signature does not yet exist, which includes zero-day exploits, newly developed malware families, and the approximately 18 percent of current malware that uses polymorphic techniques to continuously modify its own code. A polymorphic payload that changes its structure every few hours will never match any stored signature. For AI-targeted content embedded as text instructions, signature matching is essentially irrelevant: the malicious content is plain text, not executable code, and has no malware fingerprint.

Sandboxing and Behavioral Analysis

Sandboxes execute files in an isolated environment and observe their behavior. This approach is more effective against unknown threats but carries significant drawbacks. Processing time typically ranges from minutes to tens of minutes per file, making sandbox-based inspection impractical at the throughput required by email gateways or document management platforms serving large organizations. Modern malware has responded to sandboxing with sophisticated evasion techniques: virtual machine detection, timing-based dormancy that delays execution until sandbox analysis has timed out, artificially large file sizes that exceed sandbox size limits, and environment fingerprinting that prevents execution when specific analysis tools are detected.

Sandboxes also provide no protection against AI-targeted instructions embedded in files, because those instructions do not execute code and produce no behavioral signals that a sandbox would flag.

Reactive Detection and Incident Response

Many organizations still rely primarily on detecting and responding to threats after initial execution, using SIEM correlation, EDR behavioral signals, and threat hunting. This approach accepts some level of compromise as inevitable and focuses on containing damage after a breach. It does not prevent the initial compromise, and for threats that operate stealthily over

extended periods, such as advanced persistent threats or credential-stealing infostealers, significant damage may already have occurred before detection.

None of these conventional approaches adequately address the combination of threats now targeting files in AI-enabled enterprise environments. A new architectural layer is needed that operates before execution, before user interaction, and before any AI system processes the file.

How FileDNA™ Works: Structure Over Signature

FileDNA™ is built around a fundamentally different security principle. Instead of asking whether a file matches a known malware signature, or observing what the file does when executed, FileDNA™ asks a more basic question: what does this file contain, and does any of that content violate the expected structural rules for this file type? Anything that does not belong is removed. The remaining, compliant content is reassembled into a clean, functional file that is safe to open, process, and pass to an AI system.

This approach is entirely independent of any threat signature library. A zero-day exploit, a piece of polymorphic malware no security vendor has seen before, and a newly developed AI-targeted prompt injection payload are all subject to the same analysis: does this content belong in this file according to its format specification? If not, it is removed.

The Five-Stage Processing Pipeline

- **Ingest.** FileDNA™ receives the file via the RESTful API or SDK integration. Any supported format is accepted, including nested containers and multi-level archives. The file is queued for structural analysis without being executed at any stage.
- **Deconstruct.** The engine performs deep structural parsing of the entire file. Every embedded object is identified and extracted: macros, scripts, JavaScript layers, executable attachments, annotation objects, image payloads, metadata fields, linked objects, and external references. For archive types, this process recurses through unlimited nesting levels until every embedded object has been identified and mapped.
- **Analyze.** Each extracted object is evaluated against the permitted content rules for the container format and the organization's configured security policy. Proprietary detection algorithms examine structural characteristics, encoding patterns, obfuscation signatures, and behavioral indicators. Machine learning models trained in real-world malicious samples provide an additional analysis layer for identifying sophisticated or novel threats.
- **Neutralize.** Objects that violate policy are removed or replaced. Binary executable attachments are substituted with safe placeholder content that preserves the structural integrity of the container. Malicious macros, obfuscated scripts, embedded executables, suspicious external links, hidden metadata, and AI-targeted instruction payloads are removed. The file is then reconstructed from the compliant remaining content.
- **Delivery.** The sanitized file is returned to the requesting system within seconds, accompanied by a detailed report documenting every action taken, every object

removed, and the classification of each detected threat. The report is structured for integration into SIEM platforms, security dashboards, and analyst workflows.

Key Technical Capabilities

- **True Format Detection.** FileDNA™ identifies each file by examining its actual binary structure, not its filename extension. An executable renamed with a .docx extension is identified and handled as an executable. A malformed PDF designed to confuse parsers is identified by its structural characteristics. This prevents extension-spoofing attacks from bypassing analysis.
- **Unlimited Embedding Depth.** Malware authors frequently nest malicious content inside multiple levels of containers: a ZIP archive containing another archive containing a password-protected container with the payload. FileDNA™ follows every embedding level until all content has been extracted and analyzed. Depth is limited only by available system memory, not by a fixed nesting limit.
- **Polymorphism-Resistant Analysis.** Because FileDNA™ operates structurally rather than by signature matching, polymorphic malware that constantly rewrites its own code provides no evasion benefit. Every version of a polymorphic payload shares the same structural violation: it is executable content embedded where no executable content should exist. FileDNA™ removes it regardless of what specific bytes it happens to contain at analysis time.
- **AI-Targeted Content Neutralization.** FileDNA™ sanitizes the embedded content, metadata, annotation layers, linked objects, hidden text fields, and external references that represent the primary delivery mechanisms for indirect prompt injection attacks. Documents processed by FileDNA™ are safe to pass to AI assistants, coding agents, document processing pipelines, and agentic workflows because the instruction-carrying vectors have been removed.
- **Zero Code Execution.** FileDNA™ performs all analysis in a static environment. No file is ever executed during analysis. This eliminates the detonation latency of sandboxing approaches, removes the risk of zero-day sandbox escapes, and ensures that even the most dangerous content cannot cause harm during the analysis process.
- **Real-Time Processing Speed.** FileDNA™'s proprietary analysis and reconstruction engine processes most files in under a few seconds. This makes it practical for deployment at email gateways, network inspection points, and document management platforms without introducing perceptible latency into user workflows.

Supported File Formats

FileDNA™ provides comprehensive coverage across more than 50 file format families, spanning every version of the Microsoft Office suite, Adobe products, web formats, scripting languages, archive types, and image formats.

File Type	Supported Versions	File Extensions
MS Word	Word 6, 95–97, 2000–2003, 2007–2019, Office 365	.docx .docm .dotx .dotm .doc .dot .xml
MS Excel	Excel 4-5 Macro-Sheets, 95–97, 2000–2003, 2007–2021, Office 365	.xlsx .xlsm .xlsb .xltm .xlam .xls .xla .xlb .xlc .xld .xlk .xll .xlm .xlt .xlv .xlw
MS PowerPoint	PowerPoint 97–2003, 2007–2021, Office 365	.ppt .pps .pot .pptx .pptm .ppsx .ppsm .ppam .potx .potm
MS Publisher	Publisher 95–98, 2000–2003, 2007–2021, Office 365	.pub
MS Visio	Visio 2000–2003, 2007–2024	.vsd .vst .vtx .vsdx .vsdm .vssm .vstx .vstm
MS Outlook	Outlook 97–2003	.msg
Ichitaro	Structured Storage versions	.jtd .jtdc
MS WordPad	Windows 95, 2000/XP	.rtf .wri
Web Formats	All versions	.xml .htm .html .url .hta
Scripting	All versions	.js .vbs .vba .vbe .jsp .vb .py .ps1
Archives	ZIP archive content, nested levels	.zip .7z .rar .gz .tgz and more
Excel Sylk	All versions	.slk
Excel delimited	All versions	.csv
Adobe PDF	All versions	.pdf
Adobe Flash	Up to version 10	.swf
Shortcut / Link	Windows	.lnk
Command Prompt	Windows	.cmd .bat
Binary Executable	As attachments (replaced with safe content)	.exe .dll .com .sys and more
Graphic Formats	All common types	.jpg .jpeg .bmp .gif .png .ico
Embedded Scripts	Any graphics-embedded scripting	Detected within any container
Scalable Vector Graphics	All versions	.svg
More types	…	…

Note: Binary executable files detected as embedded attachments are replaced with safe, inert placeholder content that preserves the structural integrity of the container file. The container document remains usable; the threat is neutralized.

CyberQuay, Inc. © 2026 | FileDNA Whitepaper

Deployment and Integration

FileDNA™ is designed for broad compatibility and straightforward integration into existing security infrastructure. It requires no specialized hardware, no third-party runtime dependencies, and no persistent agent installation on endpoints.

API and SDK Integration

- **RESTful API.** FileDNA™ exposes a comprehensive RESTful API that allows any system capable of making HTTP requests to submit files for analysis and receive sanitized files and analysis reports in return. Integration with email security gateways, web proxies, document management systems, SIEM platforms, and AI data pipelines can typically be completed in days.
- **SDK for Developers.** A full software development kit is available for teams that prefer direct binary integration. The SDK supports embedding FileDNA™ analysis directly within existing cybersecurity products, threat intelligence platforms, and custom security tooling. Both Windows and Linux environments are supported with no third-party dependencies required.
- **Platform-Agnostic Architecture.** FileDNA™ runs on any current version of Microsoft Windows and a wide range of Linux distributions. It does not depend on external tools, virtualization environments, or cloud service dependencies. Deployment options include cloud-hosted service, on-premises installation, and hybrid configurations that meet data residency and sovereignty requirements.

Compatibility with Security Platforms

FileDNA™ is compatible with and designed to complement the full range of enterprise security platforms:

- SIEM (Security Information and Event Management): Analysis reports feed directly into SIEM correlation engines, providing structured event data on every file processed.
- EDR (Endpoint Detection and Response): FileDNA™ operates as a pre-execution layer that reduces the volume of threats reaching the endpoint, complementing EDR's post-execution behavioral detection.
- IDPS (Intrusion Detection and Prevention Systems): File-level sanitization at network ingress points reduces the malicious payload volume that IDPS engines need to process.

- NGFW (Next-Generation Firewalls): SSL inspection integration allows FileDNA™ to analyze files traversing encrypted channels.
- Email Security Gateways: All inbound attachments can be processed by FileDNA™ before delivery to user inboxes.
- Document Management Systems: FileDNA™ scans every uploaded document before storage and every downloaded document before delivery.
- AI Data Processing Pipelines: Documents and files fed to AI systems, including training data repositories, document summarization tools, and agentic workflow platforms, are sanitized before AI processing.

The Structural Gap: Why the Detection Industry Has Not Built CADR

A reasonable question follows from everything above. If prevention at the file layer works, why have the large platform vendors not simply built it? The answer is not incompetence, and it is considerably more durable than that. Six structural forces, none of which is likely to reverse in a single quarter, keep the capability outside the core offering of every major detection platform.

1. Detection Is the Business Model, Not the Outcome

An endpoint detection platform is sold on telemetry volume, alert fidelity, and analyst efficiency. A managed detection service is sold on the number of events a human team will review on the customer's behalf. An extended detection platform is sold on the breadth of signals it can correlate. Every one of those metrics is a function of how much suspicious activity flows through the system. When a disarmament layer is installed upstream, weaponized documents stop arriving, macro-laden spreadsheets stop arriving, and embedded scripts and exploit-bearing images simply stop existing by the time they reach the endpoint. The alert queue does not get better. It gets quieter. For the buyer, that is the entire point. For the vendor whose renewal conversation is built around the value of the alerts it surfaced last quarter, it is an awkward thing to sell.

2. Detection Vendors Are Architecturally in the Wrong Place

An endpoint agent lives on the endpoint. It hooks the kernel, watches process creation, monitors memory, and observes what code does once it is running. That vantage point is excellent for what it is designed to do, and it is by definition downstream of the file's arrival. Content disarmament belongs somewhere else entirely. It belongs at the mail gateway, at the web proxy, at the file upload endpoint of a customer portal, at the ingestion point of a document management system, at the boundary of a shared drive, and increasingly at the entrance to the retrieval pipeline feeding an AI model. Those are transit points rather than endpoints, and serving them requires a different deployment architecture, a different set of integration partners, and a different performance profile.

3. Reconstruction Carries a Liability Profile Detection Vendors Do Not Want

When a detection engine is wrong, the consequence is a false positive. An analyst wastes twenty minutes and the vendor tunes a rule. When a reconstruction engine is wrong, the consequence is that the finance director's quarterly model stops calculating, or the legal team's redlined contract arrives with its tracked changes silently flattened. That failure generates an executive escalation, and it lands on the vendor who touched the file. A platform vendor with a

hundred thousand enterprise customers is rationally reluctant to introduce a capability whose failure mode is corrupting the customer's own documents.

That reluctance is precisely the moat. The hard engineering problem is not identifying that a macro is dangerous, which most security vendors can do. The hard problem is rebuilding hundreds of file formats with enough fidelity that the user never notices anything happened. That is deep, unglamorous, format-by-format work accumulated over years, and it is not something a platform vendor can spin up in two quarters because the category began to look attractive.

4. The Category Was Filed Under Someone Else's Name a Decade Ago

The lineage of disarm-and-reconstruct technology runs through a set of vendors who positioned it, quite reasonably at the time, as an email security feature. It was catalogued as a gateway add-on and it appeared as a checkbox in secure email gateway comparisons. Once a capability has been filed under a category, the buying committee stops looking for it anywhere else. Meanwhile the actual threat surface moved. Files no longer arrive primarily by email. They arrive through customer upload portals, collaboration platforms, third-party integrations, partner data exchanges, supply chain document flows, and now through the ingestion pipelines feeding retrieval-augmented AI systems. The old category boundary was drawn around a channel that no longer represents most of the risk, and nobody has redrawn it.

5. The Industry's Attention Is Pointed the Other Way

Research and marketing spend is currently chasing a single narrative, which is the autonomous security operations center: the AI analyst that triages, correlates, investigates, and responds without a human in the loop. Read that narrative carefully and notice what it is. It is detection, made faster. The premise remains that threats will arrive, that alerts will fire, and that the winning move is to process them more intelligently. Prevention-first is an unfashionable argument in a market that has just spent billions funding a better response function, and it is difficult to build a compelling stage presentation around the absence of an incident.

6. The Scorecards Do Not Measure It

Vendors optimize for the tests they are graded on, and the tests in this industry are built around detection. Adversary emulation evaluations score whether a product generated telemetry for a given technique and whether it produced an alert. Independent lab tests measure detection rate and false positive rate. Analyst grids weight response capability, investigation workflow, and threat hunting depth. Every one of those frameworks presupposes a product whose job is to

notice things. A product that neutralizes a weaponized document such that the technique never executes does not score well on a matrix of techniques, because it does not participate in the matrix. On a detection scorecard, a perfect prevention outcome is indistinguishable from a product that was not installed. The measurement infrastructure of the industry is structurally incapable of rewarding the outcome the industry claims to want.

Complementary Control, not a Replacement

None of this is a criticism of detection and response. The engineering behind modern endpoint, extended, and managed detection platforms is world class, the outcomes they deliver are real and measurable, and no serious security program should attempt to operate without them. The argument is one of layering rather than replacement. Detection and response answer the question of what to do once a threat is present in the environment. Structural content disarmament answers a different and prior question, which is whether the threat needs to be present at all. FileDNA™ does not compete with EDR, XDR, or MDR. It reduces the volume of file-borne material those platforms are asked to reason about, so that their considerable analytical capability can be applied to the threats that genuinely require it.

The strategic question for a board is therefore not whether the organization's detection stack is good. It very probably is. The question is whether the organization has any control anywhere in its architecture whose success condition is that the alert never fires. If every dollar of file-borne threat spend is allocated to products that only earn their keep after the file has arrived, then the security program has a structural dependency on being fast enough, every single time, forever. That is a defensible position right up until the first time it is not.

The Economics of Prevention for Small and Mid-Sized Organizations

The enterprise security stack was never designed for small and mid-sized business economics. Layered architectures built around the major platform vendors assume an organization can afford dedicated security operations analysts, managed detection subscriptions, expensive endpoint telemetry collection, cloud log ingestion, continuous incident response operations, and extensive security engineering resources. For a company with fifty, one hundred, or even three hundred employees, that approach quickly becomes financially untenable. Attackers know this.

Real-world incident data tells a simple story about how those organizations are actually compromised. According to ESET's SMB readiness research, the most common initial vectors remain phishing emails, malicious document attachments, weaponized PDFs, infected Office documents, embedded JavaScript payloads, malicious archive containers, disguised download installers, and browser-delivered payloads hidden inside legitimate file formats. Attackers continue relying on files because files remain one of the most trusted objects inside business workflows. A finance employee opens an invoice, a sales team downloads a proposal, a procurement officer receives a spreadsheet from a vendor, and the compromise begins immediately.

Most mainstream products are reactive by design. An attachment is delivered, the user opens the file, code executes, behavior monitoring begins, detection engines attempt to identify suspicious activity, endpoint agents generate alerts, security teams investigate, and containment begins. That sequence works in an enterprise with a mature security operations center. It works far less well in an organization with no security analysts, no incident response team, no continuous monitoring, and limited forensic capability. By the time detection occurs, the compromise has already started and the business enters the expensive recovery phase.

Table 3. Typical annual cost structure of a detection-oriented security stack.

Cost Component	Typical Annual Cost
Endpoint detection platform	$15,000 to $40,000 annually
Managed detection and response service	$20,000 to $80,000 annually
SIEM or cloud log ingestion	Variable recurring cost, scales with telemetry volume
Incident response retainer	Additional annual cost
Employee downtime during incidents	Unbudgeted, incurred at time of breach
Ransomware recovery	Unbudgeted, incurred at time of breach
Compliance and insurance impact	Rising premiums and audit exposure

File-centric prevention changes this calculation by eliminating one of the most common attack vectors before downstream security systems are forced to respond. The cost of prevention is a fraction of enterprise-grade detection infrastructure, and more importantly it eliminates the hidden secondary costs that never appear on a subscription invoice, including operational downtime, business interruption, forensic investigation, customer trust damage, data recovery, ransom negotiation, and legal liability exposure. For organizations operating on tight margins, avoiding a single major incident can offset years of cybersecurity spending.

Worked Example: An Invoice Fraud Campaign

An accounts payable employee receives a vendor invoice named Invoice_June_2026_Approved.pdf. The document appears entirely legitimate. Internally, it carries an embedded phishing URL, a hidden JavaScript redirect, a credential harvesting page, and a counterfeit payment portal.

Under traditional protection, the file reaches the inbox, the user opens the document, the phishing sequence begins, and credentials are stolen. Detection, if it occurs at all, occurs after the credentials have already left the organization.

Under FileDNA™, the file enters the gateway, its internal active content is analyzed, the embedded malicious objects are removed, a sanitized PDF is reconstructed, and a safe document is delivered. The user never encounters the payload. The attack chain ends before execution, and no alert, no investigation, and no recovery is required.

Cyber resilience is conventionally described as the ability to recover quickly after an attack, but recovery should not be the primary strategy for organizations with limited resources. The more effective strategy is stopping attacks before security teams are needed at all, and that requires moving protection upstream. For organizations of this size, prevention is no longer only a matter of security architecture. It is financial strategy.

Use Cases

Email Security Gateway

Email remains the single most common initial access vector for enterprise malware. FileDNA™ deployed at the email perimeter analyzes every inbound attachment before it reaches user inboxes. Phishing payloads delivered via Office documents, PDFs, ZIP archives, and script files are neutralized at the gateway. Outbound email scanning prevents data exfiltration through weaponized documents and stops compromised internal accounts from distributing malicious content to external contacts.

AI System Data Sanitization

Organizations deploying AI assistants, coding agents, document processing tools, and agentic workflow automation face growing exposure to indirect prompt injection attacks through the files these systems process. FileDNA™ provides a sanitization layer between external file sources and AI processing pipelines, ensuring that documents, repositories, emails, and web-sourced content are stripped of AI-targeted instruction payloads before any AI system reads them. This includes removing hidden metadata, annotation-layer instructions, embedded directives in image alt text, and obfuscated text fields that are not visible to human readers but are accessible to AI models.

Document Management and File Sharing

Enterprise document management platforms, collaboration tools such as Microsoft SharePoint, and file sharing services represent high-value targets for malware distribution because they aggregate large volumes of documents from diverse sources. FileDNA™ deployed at the storage layer ensures that every file uploaded to these platforms is sanitized before it is stored, and every file retrieved is verified clean before delivery. This prevents a single compromised document from becoming a source of internal malware propagation.

Network Gateway Inspection

FileDNA™ deployed at network ingress and egress points provides inspection of all file transfers traversing the network boundary. This complements firewall and IDS capabilities by adding a content-level analysis layer that can identify and neutralize threats concealed inside legitimate-seeming file transfers. For organizations with strict data governance requirements, outbound scanning prevents unauthorized exfiltration of sensitive data embedded in documents.

Internet Service Provider Infrastructure

ISPs and content delivery networks can deploy FileDNA™ at scale to protect subscriber traffic from file-based threats before malicious content reaches the endpoint. Proactive content disarmament at the network level reduces the downstream impact of large-scale phishing campaigns and malware distribution networks, improving service quality and reducing the volume of security incidents affecting customers.

Desktop and Endpoint Protection

For individual users and small to mid-sized organizations without enterprise gateway infrastructure, FileDNA™ is available as a desktop application that provides on-demand and automated file scanning. Users can scan any file before opening it, ensuring protection even when the file arrives through channels not covered by gateway-level controls.

Security and Operational Benefits

Proactive Threat Elimination

FileDNA™ operates before malware executes, before the user opens the file, and before any AI system processes the content. This proactive model eliminates threats at their point of entry rather than attempting to contain damage after compromise has occurred. The shift from reactive detection to proactive disarmament is consistent with zero-trust principles and significantly reduces mean time to protection.

Zero-Day and Polymorphic Threat Coverage

Because FileDNA™ operates on file structure rather than malware signatures, it provides equivalent protection against known threats, zero-day exploits, and polymorphic malware that continuously rewrites its own code to evade signature matching. Security teams are no longer dependent on signature database update cycles to be protected against novel threats.

Reduced Alert Fatigue

Rather than flagging files as suspicious and queuing them for analyst review, FileDNA™ produces definitive outcomes: a file is either cleaned and returned, or rejected according to configured policy. This eliminates the large volume of uncertain verdicts that consume analyst time without producing clear security value, allowing security teams to focus on confirmed incidents and strategic security work.

Operational Continuity

The FileDNA™ processing pipeline is designed to maintain the functional integrity of every file it sanitizes. Users receive a clean version of the document they expected, not a quarantine notification or an absent file. Business workflows continue without interruption even as threats are neutralized in the background.

Compliance and Audit Readiness

Every FileDNA™ analysis action is logged in a structured audit trail documenting the file identity, analysis timestamp, content elements evaluated, threats detected, actions taken, and outcome. This audit log supports compliance with cybersecurity regulatory requirements, data governance frameworks, and cyber insurance reporting obligations.

Conclusion

The threat landscape facing organizations in 2026 has grown beyond what signature-based and reactive security tools were designed to address. Files have always been the primary delivery mechanism for malware. In AI-enabled enterprise environments, they have also become the primary delivery mechanism for a new class of attack that targets the AI systems processing those files.

FileDNA™ addresses both dimensions of this challenge with a single architectural approach: structural analysis and active-content disarmament that operates independently of signatures, produces clean verified files in seconds, and generates no false positives requiring analyst review. It protects the human recipients of documents and the AI systems that process them using the same engine and the same pipeline.

Organizations that integrate FileDNA™ into their security stack at the email gateway, document management layer, network perimeter, or AI data pipeline gain a proactive security layer that stops threats before they can execute, manipulate, or exfiltrate. In a threat environment where the file is the payload, the most effective defense is to ensure that every file arriving in your environment is, structurally and verifiably, clean.

Three conclusions follow from the evidence assembled here. First, the apparent diversity of the threat landscape is largely superficial, because the delivery layer is remarkably uniform across otherwise unrelated criminal enterprises. Second, generative AI guarantees that the number of unique malware variants, and the number of unique AI-targeted instruction payloads, will continue to grow faster than any detection catalog can track. Third, and most importantly, none of this changes the fundamental point of entry.

The file remains the front line. Securing it is the highest-leverage decision available to a security organization, and it is one of the few investments in cybersecurity whose value increases rather than decays as attacker capability improves. FileDNA™ decomposes every inbound file, removes active and anomalous content, and rebuilds a clean, fully usable version before it reaches a user, a system, or a model. There is no signature, no verdict, and no alert. The threat is not detected. It is gone.

> FileDNA™ shifts security from chasing malware variants to establishing trust at the file layer. Every file that passes through FileDNA™ arrives clean, verified, and safe for both human and AI use.

Disclaimer

The competitive and market-structure analysis presented in this white paper represents the opinion and commercial perspective of CyberQuay, Inc. It is offered as strategic commentary on the structure and incentives of the endpoint and extended detection markets, and it should be read as such rather than as a statement of fact about the intentions, internal roadmaps, or product decisions of any named company.

CyberQuay holds the detection, response, and managed security vendors referenced in this document in high regard. The engineering behind modern EDR, MDR, and XDR platforms is genuinely world class, and the outcomes those platforms deliver are real, measurable, and indispensable. Nothing here should be read as an assertion that those products are ineffective, poorly built, unnecessary, or of diminished value. Our argument is one of layering rather than replacement.

The observations regarding commercial incentives, architectural positioning, category taxonomy, and evaluation methodology are structural in nature. They describe forces that act on organizations over time rather than choices made by individuals, and they are not intended to impute bad faith or negligence to any vendor. CyberQuay, Inc. has no commercial relationship with, and no financial interest in, any of the companies named in this document. All product names, trademarks, and registered trademarks are the property of their respective owners, and their use here is for identification and commentary purposes only. This document does not constitute investment advice.

References

The sources below support the threat landscape, market structure, and economic analysis presented in this white paper. All vendor pages were accessed and verified in July 2026. Product names and capabilities reflect vendor documentation as published at that time and are subject to change.

Threat Research and Industry Data

Cybersecurity Ventures. Cybercrime To Cost The World $10.5 Trillion Annually By 2025. https://cybersecurityventures.com/cybercrime-damage-costs-10-trillion-by-2025/

IBM Security. Cost of a Data Breach Report 2024. https://www.ibm.com/reports/data-breach

Verizon. 2026 Data Breach Investigations Report. https://www.verizon.com/business/resources/reports/dbir/

ESET / WeLiveSecurity. SMB Cyber Readiness: The Road to Resilience Starts Here. https://www.welivesecurity.com/en/business-security/smb-cyber-readiness-road-resilience-starts-here/

ESET / WeLiveSecurity. Unpacking SMB Cyber-Readiness and What Makes or Breaks It. https://www.welivesecurity.com/en/business-security/smb-cyber-readiness-what-makes-breaks-it/

FortiGuard Labs. PureLogs Infostealer Campaign Analysis. https://www.fortinet.com/blog/threat-research

AI Security and Prompt Injection

OWASP. Top 10 for Large Language Model Applications, LLM01:2025 Prompt Injection. https://genai.owasp.org/llmrisk/llm01-prompt-injection/

Permiso Security. ChatGPhish: Indirect Prompt Injection in Practice. https://permiso.io/blog

Lakera AI. Prompt Injection and Zero-Click Remote Code Execution in AI Agents. https://www.lakera.ai/blog

MITRE. CVE-2025-59944, Cursor IDE Configuration File Handling. https://www.cve.org/

Endpoint and Extended Detection and Response Platforms

CrowdStrike. Falcon Insight XDR. https://www.crowdstrike.com/en-us/platform/endpoint-security/falcon-insight-xdr/

CrowdStrike. Endpoint Security Overview. https://www.crowdstrike.com/products/endpoint-security/falcon-insight-xdr/

Microsoft. What Is Microsoft Defender XDR. https://learn.microsoft.com/en-us/defender-xdr/microsoft-365-defender

Microsoft. Defender Products and Services. https://learn.microsoft.com/en-us/defender/

SentinelOne. Singularity XDR Platform. https://www.sentinelone.com/platform/singularity-xdr-protection/

Palo Alto Networks. Cortex XDR. https://www.paloaltonetworks.com/cortex/cortex-xdr

Palo Alto Networks. Cortex Endpoint Detection and Response. https://www.paloaltonetworks.com/cortex/endpoint-detection-and-response

Sophos. Endpoint Powered by Intercept X. https://www.sophos.com/en-us/products/endpoint-antivirus

Managed Detection and Response Services

CrowdStrike. Falcon Complete Next-Gen MDR. https://www.crowdstrike.com/en-us/services/falcon-complete-mdr/

Microsoft. Defender Experts for XDR. https://learn.microsoft.com/en-us/defender-xdr/dex-xdr-overview

Microsoft. Defender Experts MDR. https://learn.microsoft.com/en-us/defender-xdr/defender-experts/defender-experts-mdr-overview

Palo Alto Networks. Cortex Platform, including Unit 42 Managed Detection and Response. https://www.paloaltonetworks.com/cortex

Sophos. Managed Detection and Response. https://www.sophos.com/en-us/services/managed-detection-and-response

Category Definitions and Evaluation Frameworks

CrowdStrike. What Is XDR. Vendor explainer including the Gartner definition of extended detection and response. https://www.crowdstrike.com/en-us/cybersecurity-101/endpoint-security/extended-detection-and-response-xdr/

MITRE Engenuity. ATT&CK Evaluations. https://attackevals.mitre-engenuity.org/

MITRE. ATT&CK Knowledge Base. https://attack.mitre.org/

CyberQuay Research and Commentary

CyberQuay, Inc. The Blind Spot in Plain Sight: Why the Detection Industry Has Not Built CADR. https://cyberquay.com/the-blind-spot-in-plain-sight-why-the-detection-industry-has-not-built-cadr/

CyberQuay, Inc. Effective Prevention Against AI-Driven MaaS and RaaS Malware. https://cyberquay.com/effective-prevention-against-ai-driven-maas-and-raas-malware/

CyberQuay, Inc. SMB Cybersecurity: Why Prevention-First Technologies Like FileDNA CADR May Be the Missing Layer. https://cyberquay.com/smb-cybersecurity-why-prevention-first-technologies-like-filedna-cadr-may-be-the-missing-layer/

CyberQuay, Inc. The File Is the Payload: How Ordinary Documents Are Becoming Weapons for AI-Driven Attacks. https://aidata.services/cybersecurity/the-file-is-the-payload-how-ordinary-documents-are-becoming-weapons-for-ai-driven-attacks/

CyberQuay, Inc. How Attackers Use JavaScript and PowerShell to Steal Your Credentials. https://aidata.services/data-security/javascript_powershell-steal-your-credentials/

CyberQuay, Inc. FileDNA CADR Technology. https://filedna.io/

Copyright

Copyright © 2026 CyberQuay, Inc. All rights reserved.
No part of this publication may be reproduced, distributed, or transmitted in any form without prior written permission from the publisher, except for brief quotations used in reviews, commentary, or educational discussion.
This publication is provided for informational purposes only. It does not constitute legal, regulatory, or cybersecurity consulting advice. Organizations should evaluate their own risk profile, infrastructure, compliance obligations, and security architecture before implementing any technology or security control.

About CyberQuay

CyberQuay, Inc. develops cybersecurity technology focused on securing the file formats that organizations exchange, download, upload, analyze, and process every day.
Its FileDNA technology is built around Content Analysis, Disarm and Reconstruction. The platform analyzes supported files at the structural level, removes active or suspicious embedded content, and reconstructs clean, usable versions before those files reach users, endpoints, or downstream systems.
By shifting protection closer to the point where many attacks begin, CyberQuay helps organizations reduce file-borne malware risk, strengthen existing security controls, and prevent dangerous content from becoming an active intrusion.

Contact Information

For more information, a product demonstration, or integration support, contact CyberQuay, Inc.:

- Email: support
- Website: cyberquay.com
- AI Data Services: aidata.services/
- FileDNA CADR: filedna.io