



RESEARCH REPORT

The CyberSense Research Lab: What We're Learning About Ransomware in 2026

PUBLISHED: OCTOBER 2026 | U.S. PATENT #12248574

Findings from 1,064 ransomware strains acquired and detonated by the CyberSense Research Lab between January and June 2026.


Executive Summary

Encryption is a violent thing to do to a file. It takes structured, orderly data and turns it into noise. That noise is what made ransomware detectable: Entropy spikes. Extensions change. Thousands of files get rewritten in minutes. A ransom note arrives to make sure nobody misses the point. For most of the last decade, ransomware was loud, and defenders built an entire detection industry on the noise it made.

Across the 1,064 strains the CyberSense Research Lab detonated in the first half of 2026, a significant percentage does not produce noise—in some cases, it doesn't make a sound.

Some variants encrypt only fragments of a file. Some keep the original extension. Some leave timestamps exactly as they found them. Some skip encryption entirely and destroy the filesystem's ability to find the data instead. One variant, Encoder, rearranges the bytes inside a file without encrypting anything at all, so the entropy of the file never changes. The typical scans many tools have implemented would come back clean. Yet the file is destroyed.

FIGURE 1 – 1H 2026 AT A GLANCE



PART ONE

How Ransomware Is Behaving Now

The findings below act as a sort of sequence. They show how the signals data scanning tools have relied on for years are being removed one at a time, until almost nothing is left to detect.

01 Two Layers of Stealth

Ransomware hides on two layers. The first is the variant itself. The second is the damage it leaves behind in the data. They are different problems, and they elude different defenses.

02 What the Lab Measured in 1H 2026

Across the 1,064 strains studied, the Lab tracked six corruption and destruction behaviors. A single strain can show more than one, so the categories overlap. This represents analysis of the variants existing in the wild and not necessarily those used in actual cyber attacks.

FIGURE 3 – CORRUPTION AND DESTRUCTION BEHAVIORS, 1,064 STRAINS

BEHAVIOR		SHARE OF STRAINS	STRAINS
Directory-entry destruction		47.8%	509
Partial encryption		31.4%	334
Full encryption		18.3%	195
Legitimate-looking / disguised extensions		13.9%	148
Wiper behavior		7.2%	77
Faked timestamps		3.1%	33

03 Five Techniques Worth Planning For

Detection watches for side effects. Renamed extensions. Entropy spikes. Bursts of file changes. Altered timestamps. A ransom note. Each of the five samples below, all detonated and documented by the Lab in 2026, removes one of those side effects. Read in order, they take the effectiveness of common data scanning tools instruments away one at a time. They are individual samples, not measures of how widespread each technique is. But we expect their use to grow as their effectiveness becomes clear.

CyberSense's response to each of these is the same. It inspects the contents of a file and analyzes how the data changes over time at the byte level, catching subtle changes such as partial-based encryption that threshold or compression rate analysis alone would miss. Here are some examples of variants that will be challenging to detect by most scanning tools.

FIGURE 4 – FIVE TECHNIQUES, FIVE SUPPRESSED SIGNALS

VARIANT	DISCOVERED	TECHNIQUE
CryptoLocker	Apr 27, 2026	Preserve the extension
Satan	Feb 11, 2026	Low-entropy corruption
TimeTime	1H 2026	Encrypt slowly
2c3s6aax6x-enc	Jul 1, 2026	Preserve timestamps
AbyssLocker	Jan 2, 2026	Target selected sections

Preserving timestamps. 2c3s6aax6x-enc, discovered July 1, 2026, appends a “.enc” extension and allocates a new inode, while preserving mtime and ctime. Only atime updates. The analyzed sample did not encrypt file contents at all, so the corruption here is closer to disguise than payload. It drops a ransom note. Timestamp-based comparisons rely on change being visible in time, and this sample removes that signal.

Targeting selected sections of a file. AbyssLocker, discovered January 2, 2026, encrypts only a few 1 MiB sections of each targeted file and appends an 88-byte trailer. It preserves the original inode, updates timestamps, appends a randomly generated extension, and drops a ransom note. A database, a VM image, or an application file does not need full encryption to become unusable. Corrupting a small, critical portion accomplishes the same result with a far smaller footprint.

Each of these variants has a tell. CryptoLocker spikes entropy. Satan updates timestamps. AbyssLocker sends a note. Every one of them is visible to a tool looking for the right thing. That is what makes the next sample different.

04 Spotlight: Encoder

DISCOVERED: APRIL 29, 2026


Encoder is where that sequence runs out of road. Rather than encrypting a file's contents, which raises entropy and produces a reliable change signal, Encoder rearranges the existing bytes in place. The distribution of those bytes barely moves, so entropy stays close to where it started. There is no spike, no renamed file, no altered timestamp, no note. Nothing about the file announces what happened to it. The data inside is destroyed.

05 Attack Speed and Blast Radius

In the Lab's detonations, damage accumulated fast enough to outrun most response plans. Those detonations show a median attack velocity of roughly 97,321 files corrupted per hour. At that rate, an attack reaches 10,000 files in about six minutes, often before a human has finished confirming something is wrong. The median incident in the Lab's data damages roughly 16,980 files.

Speed is what turns a contained incident into a costly one. The interval between the first corrupted file and the first human response determines how much of an organization's data ends up in question. These figures come from controlled detonations, so they describe how fast these samples worked in the Lab rather than a benchmark for attacks in the field. The Research Lab plans to track this more closely in future reports, including median files affected, corruption speed, how the fastest attacks compare to typical ones, and how slow, deliberately throttled attacks differ from both. The goal is to give organizations a clearer picture of how much time they have, so recovery planning gets built around real numbers instead of assumptions.

FIGURE 6 – ATTACK SPEED AND BLAST RADIUS



WHY THIS MATTERS

Six minutes is less time than most escalation paths

06 The AI Question

Every conversation about ransomware in 2026 assumes AI is behind it. The Lab detonated 1,064 strains looking for that at the moment of impact, but did not find it.

That is a narrow claim, and it is worth stating precisely. The Lab sees the final, destructive stage of an attack: the executable, its behavior, and the file changes it leaves behind. Proving that a frontier model wrote or steered a binary would take its development history and the full attack workflow, and a detonation shows neither. What the Lab can say is that none of the samples it ran did the one thing an AI-driven payload would be expected to do, which is make choices about the data in front of it.

AI's footprint in 2026 sits earlier in the attack lifecycle: reconnaissance, vulnerability discovery, initial access, target discovery, lateral movement, command and control. By the time the payload runs, the decisions have already been made. The Lab is watching for the shift that would change that, and it knows what it would look like. Dynamic selection of high-value files. Different corruption methods applied to different files in the same run. Adaptation based on what a file contains or which application owns it. None of it has turned up yet. These are hypotheses the Lab is testing against, not findings it is reporting.

That pattern lines up with what the broader industry is seeing. Palo Alto Networks' 2026 Global Incident Response Report found that AI is changing the speed and scale of cyberattacks more than it is changing the attacks themselves, with threat actors compressing reconnaissance and development work that once took days into a matter of hours. ReliaQuest's 2026 research found lateral movement happening in as little as four minutes, down from hours just two years earlier, and estimated that roughly 80% of the ransomware groups it tracked are using AI or automation somewhere in their operations. Both point the same direction as the Lab's own detonations. AI is compressing everything that happens before impact. The corruption itself, the behavior documented throughout this report, is still the work of the

PART TWO

How the Research Lab Tracks It

The findings in Part One came out of controlled detonations. Part Two describes how those detonations happen, what they prove, and why ransomware recovery requires deep content data integrity validation.

07 The Daily Funnel

The Lab's intake is automated and runs without pause, covered by U.S. Patent #12248574, "AI Identification of Computer Resources Subjected to a Ransomware Attack." New variants are filtered and detonated as they appear.

The narrowing is the point. Volume arrives from malware aggregation and research sources, including VirusTotal, ReversingLabs, and academic feeds, and most of it is redundant. A single family can generate thousands of near-identical samples. Automated and AI-driven selection strips that down to what looks genuinely new. Detonation costs time and analyst attention, so only what survives selection gets run, and only what behaves unusually gets documented in full.

Over the first half of 2026, that produced the 1,064 distinct strains this report covers, the newest addition to a cumulative total of more than 7,500 variants analyzed since the Lab began.

FIGURE 7 – THE DAILY FUNNEL

STAGE	VOLUME	WHAT HAPPENS
1. Acquire	~3,500 / day	Samples arrive from aggregation and research sources
2. Select	~200 /	

08 Why the Research Lab Exists

Four things define the Lab's approach:

- It detonates real ransomware in an isolated environment, so nothing is simulated or inferred.
- It observes damage at the file and content level, capturing what the malware did to the data.
- It converts new behaviors into structured research and model training data.
- It studies emerging techniques to future-proof its testing methods against potential future variant behavior.

The Research Lab's objective is to discover new behaviors, which is a different question from how common any one of them is. Section 9 sets out what that distinction rules in and out.

A second process turns that output into a validated detection model, in four stages. Real-world data collection draws on millions of anonymized customer data samples and file operations. Live ransomware detonation is the funnel described above. Automated dataset generation turns what the Lab observes into millions of virtual protected data sets and simulated attack scenarios, classified by behavioral compromise indicators. Rigorous testing and validation refines the model against real-world and simulated samples to hold it to its detection SLA of 99.99%.

The two pipelines feed each other. The funnel finds new behavior. The four-stage process turns it into detection.

WHY THIS MATTERS

Recovery comes down to choosing a copy of your data to restore from and trusting it is clean. That decision is only as good as the tool validating it, and that tool has to recognize corruption built to evade detection.

09 Methodology and Limitations

</

PART THREE

How This Shapes CyberSense

The Lab's findings drive innovation and ensure customers can recover from even the most sophisticated ransomware attacks.

10 From Detection to Development

Each behavior documented in Part One becomes an input to CyberSense's detection model, by way of the Lab's four-stage training process described in Part Two. A technique like Encoder's entropy preservation does not stay a research curiosity. It becomes a training case that pushes the model to evaluate file structure and internal consistency alongside entropy, rather than relying on entropy as a primary signal. A technique like AbyssLocker's partial encryption pushes the model toward evaluating smaller sections of a file rather than the file as a whole. This is what the 64.7% measure from Section 1 is for: tuning detection rules against exactly these kinds of variation.

CyberSense organizes what it has learned around four connected jobs. It validates that mission-critical data stays clean and recoverable on an ongoing basis. It detects ransomware-induced corruption—even advanced stealth approaches of corruption, as it happens. It responds with a forensic account of what was affected. It recovers by pinpointing clean data with confidence. The first two have to hold before the others mean anything. A forensic account and a clean recovery point are only ever as good as the detection that produced them.

11 What this Means for Recovery

Security tooling upstream, including malware identification, suspicious-process detection, network monitoring, and encryption-activity alerts, answers one question: did something unusual happen? Once an organization reaches the

Conclusion

In the Lab's IH 2026 sample set, ransomware changed on two fronts, and both land on the same recovery problem.

At the executable level, polymorphic variants regenerate their signature with every infection. A fingerprint taken from one victim will not match the same ransomware on the next, which is why static identification now carries less of the detection load than it used to.

At the data level, a growing share of samples suppressed the signals common data scanning tools have relied on for years. Partial encryption limits the footprint. Extensions stay legitimate. Timestamps stay untouched. Entropy stays flat. Taken together, they describe ransomware that has learned to leave no mark on the surface of a file.

The Research Lab exists to find techniques like these early, and its four-stage process turns each one into a stronger CyberSense model. Data that looks unchanged is not necessarily clean data. Understanding what happened inside the data, and building that understanding directly into CyberSense, gives customers a critical last layer of defense. It elevates cyber resilience overall and turns recovery from a ransomware attack into a proven, predictable process.

// Encoder showed us that a file can look untouched and still be compromised. That's why we read the data instead of trusting surface scans.

CYBERSENSE RESEARCH LAB

About Index Engines

At Index Engines, we are experts in Cyber Resiliency, helping organizations build an infrastructure where trusted data is available and reliable. Our leading solution, CyberSense, provides a 99.99% SLA for detecting ransomware corruption. CyberSense empowers organizations to confidently navigate cyber challenges, mitigate risks, and quickly recover to normal business operations in the ever-evolving cyber landscape.