INDEX ENGINES™

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.

None of this is accidental. It is how ransomware responds to detection methods that were working, and it is the thread running through this report.

We are telling a story in three parts: how ransomware is behaving now, how the Research Lab tracks that behavior, and how the Lab's findings shape CyberSense.

The findings here underscore the need for data integrity validation. If corruption can hide from the tools most commonly used, the integrity of data is not certain until it is proven to be clean.

The CyberSense Research Lab operates under U.S. Patent #12248574, covering an automated process for collecting, detonating, and analyzing emerging ransomware to train CyberSense. That process has analyzed more than 7,500 ransomware variants to date, drawing on 120 million-plus customer datasets and 300 million-plus file operations. CyberSense's detection model has been validated against 125,000 real-world samples to a 99.99% detection SLA, audited by Enterprise Strategy Group (now Omdia). This report covers six months of that ongoing research.

FIGURE 1 – 1H 2026 AT A GLANCE

1,064	**64.7%**	**39%**	**31.4%**
Ransomware strains studied in 1H 2026	Regenerated their file signature between builds	Showed directory-entry corruption	Showed partial encryption

Directory corruption: samples that left directory entries replaced, deleted, or inaccessible after the attack. Includes directory entry replacement (28%), wipers that delete files (5%), and attacks on the disk or partition table (6%). Figures include attacks that affect some or all directories.

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.

The executable layer. Most security tools identify malicious software by its digital signature, a fingerprint calculated from the contents of the file. Change a single byte and the fingerprint changes completely. The model works by capturing a known malicious file, generating its signature, and distributing that signature to scanners in the field, where it waits for a match.

Polymorphism breaks that model. A polymorphic program rewrites its own code, its packaging, or its encryption keys with every new infection. The disguise is new each time. The payload is not. A signature harvested from one infected machine will not match the same ransomware on the next one, so every new infection arrives as something the scanner has never seen. Antivirus, endpoint protection, and every other tool that leans on digital signatures loses ground against it.

The Lab measures this by looking at what changes between builds. In 64.7% of the samples studied, the code section of the executable stayed the same while its data section changed. That is signature regeneration at work: the functional payload holds steady while the fingerprint of the file morphs.

The data layer. Partial encryption, preserved extensions, preserved or faked timestamps, low-entropy corruption, throttled encryption speed, and selective targeting all shrink the visible footprint an attack leaves on the files it touches.

Entropy deserves attention here, because so much detection depends on it. Entropy measures the randomness of a file's contents. Ordinary files have structure and score low. Encryption destroys that structure and drives the score up, and that jump has been one of the most dependable ways to spot ransomware. Corruption that maintains expected entropy levels removes the signal entirely.

FIGURE 2 – TWO LAYERS OF STEALTH

Executable stealth	Data-corruption stealth
• Unique file signature on every build	• Partial and low-entropy encryption
• Code and data repackaged per sample	• Preserved timestamps and extensions
• Static fingerprints lose ground	• Directory-level corruption

© 2026 All rights reserved. Index Engines Inc. CyberSense is a registered trademark of Index Engines Inc. All products mentioned are trademarked by their respective organizations.

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 corruption	**39%**	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

Categories overlap: a single strain can show more than one behavior, so the values do not sum to 100%.
Ranked by share of the 1,064 strains studied.

A directory entry is the filesystem's record of a file, Microsoft NTFS for example. It holds the name, the location on disk, and the pointers that let the system find the data. Directory-entry corruption attacks that record. The file contents may sit untouched on disk, but the system can no longer locate them, so the data is gone for practical purposes.

Directory-entry corruption was the most common behavior in the Lab's sample set at 39%, well ahead of full encryption at 18.3%. The ordering points to techniques aimed at the filesystem's ability to find and access files. Think of it as destroying the library catalog instead of scrambling the words inside any one book. Recovery planning built around the question "does this file look encrypted" covers less ground than it once did. Directory-entry corruption is not a more sophisticated form of encryption. It is a way of skipping encryption altogether, and it is now the most common occurrence found in the Lab's analysis. What follows are five more ways the samples detonated in the Lab try to evade detection.

WHY THIS MATTERS

Recovery planning organized around encryption signatures misses the most common behavior in the Lab's set. Directory-level damage leaves file contents intact, so a check that looks for scrambled data finds nothing wrong with them.

© 2026 All rights reserved. Index Engines Inc. CyberSense is a registered trademark of Index Engines Inc.
All products mentioned are trademarked by their respective organizations.

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 the extension. CryptoLocker, discovered April 27, 2026, fully encrypts every targeted file. It preserves the original filename and inode and appends a 16-byte trailer. Timestamps update to the attack time. No ransom note follows. Detection that leans on renamed files or suspicious extensions has few clues to work with here.

Producing low-entropy corruption. Satan, discovered February 11, 2026, fully XOR-encrypts each file using an 80-byte key template that varies across the file, while preserving filename and inode. Timestamps update to the attack time. No ransom note follows. Entropy- and compression-based detection depends on a large statistical shift, and this corruption does not produce one.

Encrypting slowly. TimeTime uses a simple substitution cipher, incrementing each byte value by one, and preserves file size and inode. In sandbox testing, the attack was still running after 44 hours, having reached roughly 10,000 files. It appends a ".timetime" extension and drops a ransom note. Change rate-based monitoring depends on a burst of activity, and a deliberately slow attack stays under that threshold.

© 2026 All rights reserved. Index Engines Inc. CyberSense is a registered trademark of Index Engines Inc.
All products mentioned are trademarked by their respective organizations.

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.

Encoder preserved the filename, the original inode, the file size, the modification time, the change time, and the entropy. It rearranged the existing bytes, changed the internal organization of the file, and destroyed its usability. The access time updated during the attack, since reading the file was part of the process. No ransom note followed.

CyberSense identified the corruption by analyzing the structure and content of the file directly. That is the same principle behind every finding in this report: deep analysis uncovers what metadata or surface scans can't be relied on to find. Data integrity products that rely on high-level surface scans will report a file like this as clean. That gives an organization false confidence in data it cannot actually restore, and it turns detection and recovery into a long manual exercise.

FIGURE 5 — ENCODER: PRESERVED VERSUS CHANGED

What Encoder preserved	What Encoder changed
• Filename • Original inode • File size • Modification time • Change time • Entropy	• Internal file structure • Data usability

© 2026 All rights reserved. Index Engines Inc. CyberSense is a registered trademark of Index Engines Inc. All products mentioned are trademarked by their respective organizations.

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

97,321
Files corrupted per hour, median

6 minutes
To reach 10,000 files

16,980
Files damaged, median incident

WHY THIS MATTERS

Six minutes is less time than most escalation paths take. An alert fires. A SOC analyst picks it up. An incident responder separates it from the routine noise of backup jobs and indexing runs. A CISO communicates that production should be taken offline.

In the Lab's detonations, 10,000 files were corrupted before that sequence typically finishes.



© 2026 All rights reserved. Index Engines Inc. CyberSense is a registered trademark of Index Engines Inc. All products mentioned are trademarked by their respective organizations.

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 ransomware.

WHY THIS MATTERS

AI is shortening the runway in front of an attack, not changing what happens at the end of it. The corruption behaviors in this report are the ones your validation has to catch, and they have not been replaced by something new. What has changed is how little time you get before they start.



© 2026 All rights reserved. Index Engines Inc. CyberSense is a registered trademark of Index Engines Inc. All products mentioned are trademarked by their respective organizations.

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 / day	Representative candidates chosen for uniqueness and novelty
3. Detonate	10-15 / day	Run live and analyzed in the isolated lab environment
4. Document	1,064 / 1H 2026	Distinct strains with notable behavior recorded

© 2026 All rights reserved. Index Engines Inc. CyberSense is a registered trademark of Index Engines Inc. All products mentioned are trademarked by their respective organizations.

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

This report reflects the Research Lab's own findings, not a global attack-frequency estimate. The Lab's acquisition sources shape which samples are available, and some ransomware families generate far more samples than others. A technique's frequency in the research feed is not the same as its frequency in real-world attacks. The percentages in Section 2 also do not sum to 100%, because a single strain often shows more than one behavior. A sample that both destroys directory entries and disguises its file extension gets counted in both categories, which is why these numbers describe overlapping patterns rather than a strict breakdown of all 1,064 strains. Trend comparisons across reporting periods hold only as long as the acquisition and filtering methodology stays consistent.

© 2026 All rights reserved. Index Engines Inc. CyberSense is a registered trademark of Index Engines Inc. All products mentioned are trademarked by their respective organizations.

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.

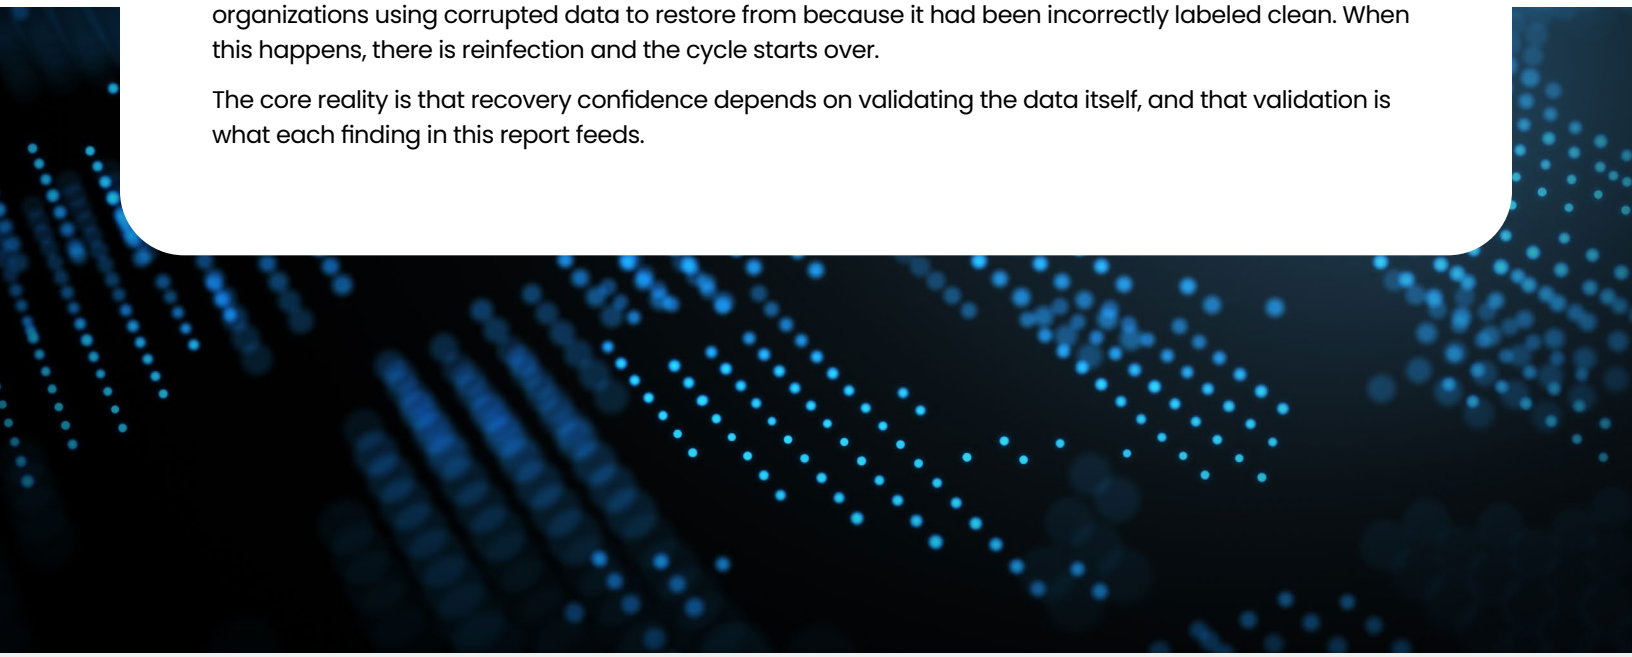
CyberSense organizes what it has learned around four connected jobs. It validates that mission-critical data stays clean and recoverable on an ongoing basis. It detects ransomware-induced corruption—even advanced stealth approaches of corruption, as it happens. It responds with a forensic account of 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 recovery stage, the question changes. Can this copy of the data be trusted?

Encoder answers that question by example. A file can carry the same filename, the same size, the same inode, the same modification time, the same change time, and the same entropy, and still be corrupted. The absence of visible change is not evidence of data integrity. What's more, common scanning tools often provide a false sense of confidence that data is clean. We have seen many unfortunate examples of organizations using corrupted data to restore from because it had been incorrectly labeled clean. When this happens, there is reinfection and the cycle starts over.

The core reality is that recovery confidence depends on validating the data itself, and that validation is what each finding in this report feeds.



© 2026 All rights reserved. Index Engines Inc. CyberSense is a registered trademark of Index Engines Inc. All products mentioned are trademarked by their respective organizations.

Conclusion

In the Lab's 1H 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.

Learn More | Contact Index Engines

© 2026 All rights reserved. Index Engines Inc. CyberSense is a registered trademark of Index Engines Inc. All products mentioned are trademarked by their respective organizations.