STAIRWELL

See the Unseen:

HIDDEN MALWARE

July 2026

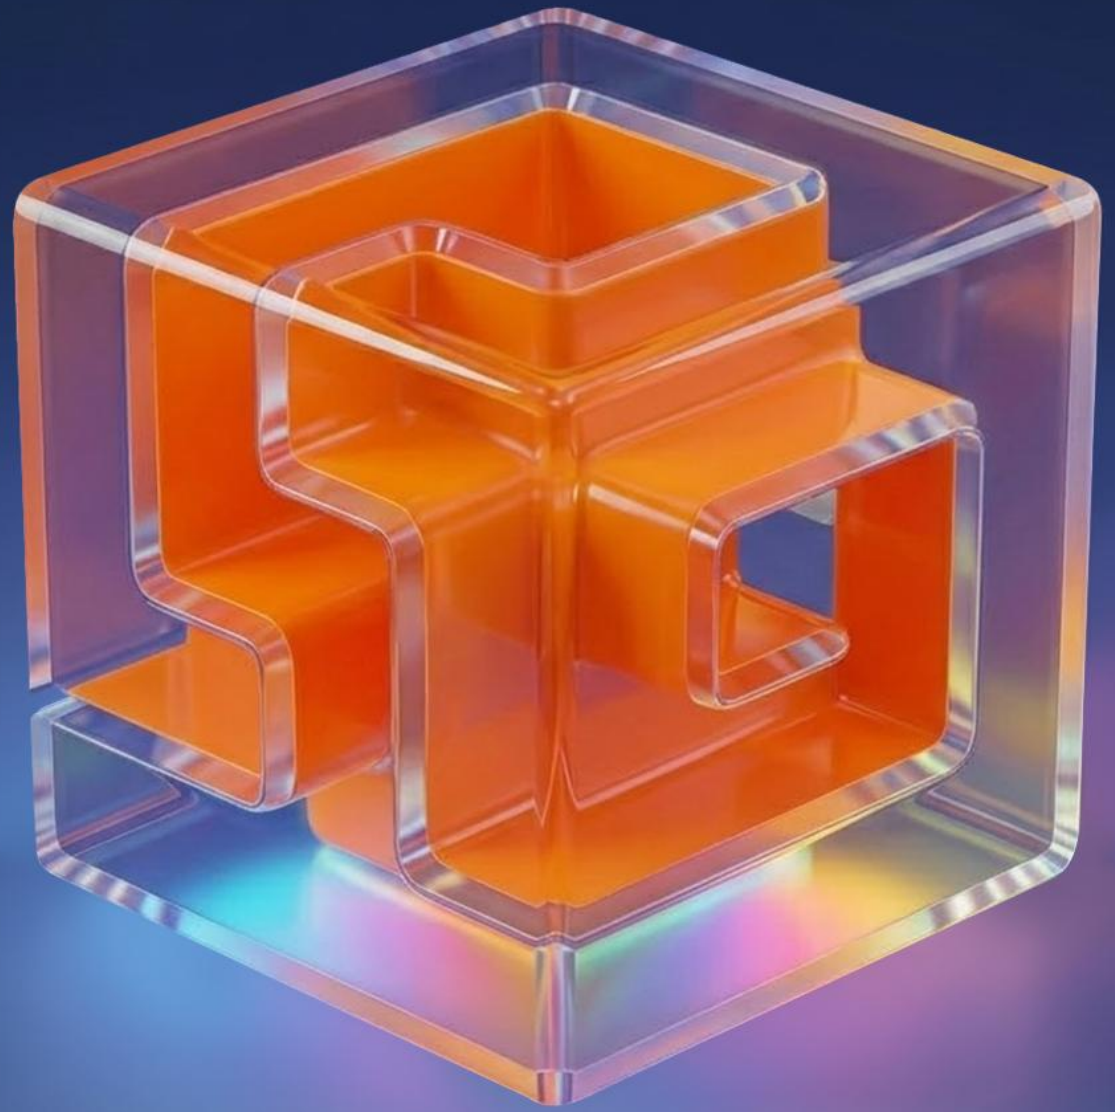

STAIRWELL.COM

Executive Summary

Stairwell analyzed 1,085 public threat reports published between March 2023 and June 2026, containing 19,418 malware hashes. Using Stairwell's Variant Discovery, we analyzed every published hash to identify related malicious files the reports never named.

KEY FINDINGS

- **74% of reports** left behind at least one malicious variant that a hash blocklist alone would have missed. Said differently, nearly **three out of every four threat reports** were incomplete if you relied only on their published hashes.
- Banking-related hashes showed the largest expansion, with **7.2x** more variants than were published.
- Ransomware expanded **4.6x**, reflecting how configurable ransomware builders continuously generate fresh, hash-unique binaries from the same underlying codebase.
- Nation-state activity revealed two different patterns. **North Korea (4.7x)** and **Russia (3.9x)** rapidly generate new variants that outpace blocklists, while **China (1.2x)** favors smaller numbers of targeted implants that often never appear on public blocklists.

CONCLUSION

The common thread across every category we analyzed is code reuse. Attackers rarely start from scratch.

AI amplifies this pattern. Rather than fundamentally changing how malware is created, it dramatically lowers the cost and time required to modify existing code at scale.

THREATS REPORTS PUBLICATIONS

From March 2023 to June 2026

1,085
Public Threats Reports

19,418
Malware Hashes

46,594
Additional Malicious Files

240%
Expansion (2.4x)

BANKING
RELATED HASHES
7.2x
MORE VARIANTS THAN PUBLISHED

Research Overview

Every week, leading cybersecurity organizations publish threat reports documenting new malware campaigns, adversary tradecraft, and indicators of compromise. This report analyzes **1,085 public threat reports** published between March 2023 and June 2026, containing **19,418 malware hashes**, to understand how much malicious activity exists beyond the indicators that are publicly disclosed.

The data was collected from the Stairwell platform in June 2026. We analyzed every report that listed one or more file hashes and ran Stairwell's Variant Discovery on a per-report basis to identify all unique malicious variants related to the published hashes.

Together, these reports provide a quantitative view of how much malicious activity remains hidden beyond published indicators.



Understanding Variant Discovery

A malware variant is a file that shares the same underlying intent and behavior as another sample but has a different hash.

Rather than developing entirely new malware, attackers typically modify payloads that already work by repacking binaries, re-signing executables, swapping strings, recompiling code, or re-encrypting payloads. These changes create new hashes while preserving the same underlying behavior.

Stairwell's Variant Discovery identifies these relationships through recursive static and behavioral analysis, shared signatures, import-table similarity, and code reuse, replacing manual reverse engineering with analysis that scales.

FINDINGS

240%

Stairwell's Variant Discovery expanded detection 2.4x in the time period reviewed



THE HIDDEN MAJORITY

71%

of malware variants never appear in threat reports

- 19,418 Hashes in Published Reports
- 46,597 Variants Detected by Stairwell

Malware continues to evolve long after a threat report is published, and Stairwell's Variant Discovery continues to identify related malicious files as they emerge. A report published in 2023 has had nearly three years for related samples to surface, while reports published more recently have had only months. **As a result, expansion ratios for newer reports should be viewed as a floor rather than a ceiling.**

But that timeline is shrinking. Attackers rarely build entirely new malware. They reuse code that already works, making small modifications that are faster and cheaper than starting over. **AI is accelerating this trend**, enabling attackers to produce more malicious variants in less time. As a result, related malware appears sooner and in greater numbers after the original sample is identified.

The Hidden Malware Gap, Year by Year



Jul 2023 — Jun 2024

Published: 6,517 | Additional: 17,650 | 2.71x Expansion

Jul 2024 — Jun 2025

Published: 6,786 | Additional: 17,322 | 2.55x Expansion

Jul 2025 — Jun 2026

Published: 4,127 | Additional: 8,195 | 1.99x Expansion

Expansion by Malware Type

Malware is always written with an objective in mind, and that objective influences how it is developed, reused, and deployed. Those differences are reflected in the expansion ratios we observed across malware types.

RANSOMWARE-RELATED HASHES (4.6x)

That is consistent with the ransomware business model, where affiliates maximize return on investment by repeatedly reusing payloads that are already known to work. Configurable builders generate fresh, hash-unique binaries while preserving the same underlying capability.

BOTNETS (3.0x) AND IoT MALWARE (3.8x)

Both rely heavily on shared or leaked codebases that are repeatedly recompiled with only minimal modifications. Because they often target devices with historically thinner security coverage, they remain effective longer, increasing the return on that code reuse.

Expansion by Malware Type

MALWARE TYPE	EXPANSION RATE	TREND ANALYSIS
RANSOMWARE	4.6 x	Configurable builders generate fresh binaries consistently.
IoT MALWARE	3.8 x	Relies on shared/leaked codebases recompiled frequently.
BOTNET	3.0 x	Minimal modifications applied to known successful payloads.
RAT / BACKDOOR	2.8 x	Targeted but iterated heavily across diverse campaigns.

Expansion by Actors and Sectors

EXPANSION BY NATION-STATE ACTOR

Nation-state actors also show distinct expansion patterns that reflect different operational models.

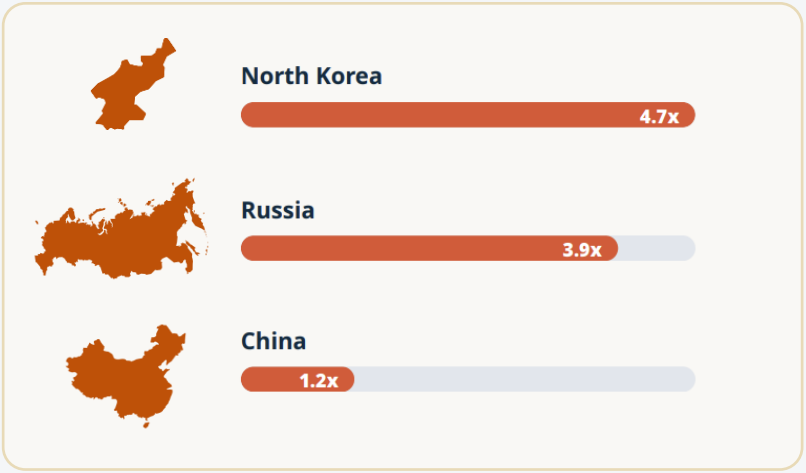

NORTH KOREA (4.7x) AND RUSSIA (3.9x)

Showed the highest expansion ratios, consistent with malware that is repeatedly recompiled, redeployed, and reused across campaigns.

CHINA (1.2x)

Showed far lower expansion, consistent with more targeted malware deployed in smaller numbers.

Together, these approaches expose two different limitations of hash-based detection: high-volume actors generate new variants faster than blocklists can keep up, while highly targeted malware often never appears on those blocklists in the first place.

EXPANSION BY SECTOR

Expansion patterns also vary by industry, reflecting the types of malware most commonly used against each sector.

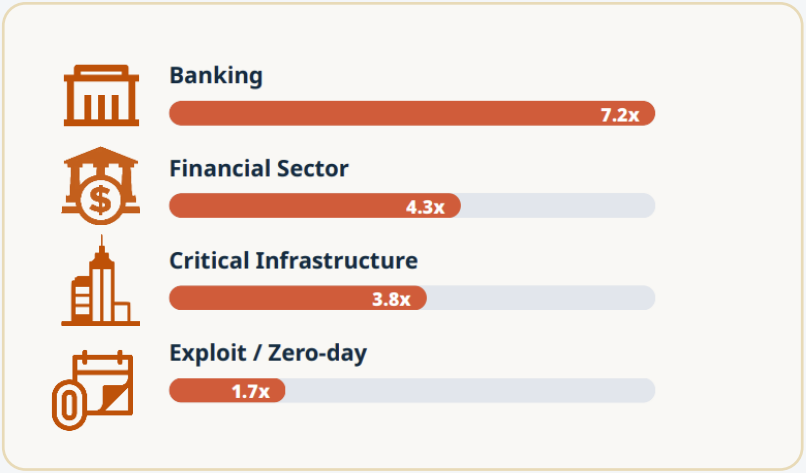

BANKING-RELATED HASHES SHOWED THE HIGHEST EXPANSION, AT 7.2x.

On average, each published banking-related hash corresponded to more than seven additional malicious files, well above the **2.4x** average across the full dataset.

FINANCIAL SECTOR-RELATED HASHES (4.3x) + CRITICAL INFRASTRUCTURE (3.8x)

followed the same pattern, reflecting sustained reuse of the same malware across multiple campaigns.

At the other end of the spectrum, **exploit and zero-day-related hashes (1.7x)** expanded far less. These reports often focus on bespoke exploits or proof-of-concept code used in a limited number of operations, leaving less opportunity for widespread variant generation.

Conclusion and Recommendations

Changing a hash is cheap, so attackers do it constantly. They repack, re-sign, swap a few strings, or re-encrypt the payload. The code keeps working, the objective stays the same, but the resulting file has a new hash that has never been published before.

Threat reports and published indicators remain valuable, but they represent a point in time. Malware continues to evolve long after a report is published, and AI is accelerating that cycle by making it faster and cheaper to generate new variants from malware that already works.

Keeping pace means treating published indicators as the beginning of an investigation, not the end. That requires:

Treating published IOCs as seeds
Hunting for related malware rather than just the exact hashes listed in a report.

Using YARA to identify related malware
Writing rules that target structure and behavior rather than exact indicators.

Continuously re-evaluating historical files
as new intelligence becomes available.

Doing that manually requires collecting files across the enterprise, preserving them over time, and repeatedly applying new intelligence to your historical data. **Stairwell automates that process, making hidden malware visible every day.**

BOOK A DEMO

See what's hiding in your environment.

ABOUT

STAIRWELL

Stairwell is the AI SOC built on ground truth. It reasons through every file that has touched an enterprise environment to find threats others miss, investigate them end-to-end, close the loop, and prove what has been contained or remediated. Stairwell preserves file history over time, enabling organizations to re-evaluate their entire environment as new intelligence arrives.

Founded by Mike Wiacek, founder of Google's Threat Analysis Group and Alphabet's Chronicle, Stairwell is backed by Sequoia, Accel, and S32, with enterprise customers across financial services, healthcare, fintech, AI, media, and gaming.

Learn more at

STAIRWELL.COM