Malware and Sha1-Hulud, TeamPCP is increasing Phoenix rebases malware Blue Shield endpoint agent against dev malware

Malware now moves faster than advisories, targets AI agents writing your code, Blue Shield blocks malicious packages and skills at the agent, endpoint, pipeline

Phoenix Security Software Supply Chain Analysis 2026 Report

LONDON, UNITED KINGDOM, June 29, 2026 /EINPresswire.com/ -- Phoenix Security today launched Blue Shield, a supply chain firewall built for a threat that has outrun the defenses most teams still rely on. Over the past year, software supply chain attacks have shifted from isolated incidents to sustained, self-propagating campaigns that arrive without a CVE, move faster than any advisory, and increasingly target the AI coding agents now writing a large share of production code. Blue Shield stops malicious packages and agent skills before they run, across the agent session, the developer workstation, and CI/CD. A free core tier is open to everyone today.

The problem: the supply chain is under sustained, accelerating attack
Phoenix release the Supply Chain Attacks Report, an analysis of the last 3 years of attacks

September 2025 was the inflection point. The Shai-Hulud worm became the first self-replicating npm malware, and Palo Alto's Unit 42 has since tracked a steady acceleration in both the frequency and technical depth of supply chain compromises, describing the registry as a force multiplier for malware distribution. What was once a nuisance is now a high-consequence threat landscape.
The pace since has been relentless, and each wave is documented by multiple independent vendors:
On May 19, 2026, a Mini Shai-Hulud wave attributed to the group TeamPCP published more than 300 malicious versions across 323 packages in a roughly 22-minute automated burst, targeting

Alibaba's AntV ecosystem and packages that represent around 16 million weekly downloads (Snyk, StepSecurity).
By early June, a descendant named Miasma had infected at least 57 npm packages and over 300 malicious versions, scanning each victim for cloud credentials and using them to spread further.

On June 1, the same family backdoored at least 32 packages in Red Hat's @redhat-cloud-services npm namespace, bypassing code review entirely (Unit 42).
In mid-May, after the Trivy scanner incident, TeamPCP open-sourced the worm, and copycats appeared within days, making attribution harder and the volume worse

Phoenix Security Blue Shield

GitHub's own ecosystem data tells the same story from another angle: npm malware advisories rose 69% year over year in 2025, the highest volume since malware tracking began (reported via Resilient Cyber). Phoenix Security's own intelligence corpus tracks 59 campaigns and 657 indexed malicious package versions from June 2024 through June 2026, with the first half of 2026 alone carrying roughly 4.5 times the malicious package volume of all of 2025.

> The attacks now move faster than the advisories meant to warn you and they have moved onto the agents writing our code. We built Blue Shield protecting agents and endpoints and released it today FREE”
>
> *Francesco Cipollone*

The attack surface moved to the AI agent

This is the part most tools cannot see. Attackers no longer limit themselves to packages. They target the developer's IDE extensions and the skills and configuration files that AI coding agents load and execute on their own. A poisoned VS Code extension, a malicious CLAUDE.md entry, a rogue MCP tool — each turns the developer's own assistant into the thing exfiltrating secrets, without a single suspicious binary ever running.
Phoenix's scanning bears this out: agent skills carry a markedly higher risk rate than IDE extensions, and more than one in four deep-scanned skills triggered a critical-risk finding. With AI now writing an estimated 30% to 41% of code at the largest engineering organizations, the agent is both the fastest-growing producer of code and the fastest-growing attack surface.

The system built to warn you has structurally fallen behind

Even when a flaw does merit a CVE, the advisory pipeline can no longer keep pace. In April 2026, NIST formally narrowed how it enriches the National Vulnerability Database, moving everything published before March 1, 2026, into a lowest-priority category it may never fully process. It cited a 263% rise in CVE submissions between 2020 and 2025 and enriched nearly 42,000 CVEs in 2025 — 45% more than any prior year — while still falling behind (NIST). A May 2026 federal audit put the unprocessed backlog above 27,000 vulnerabilities, and FIRST has forecast that 2026 will be the first year to cross 50,000 published CVEs.

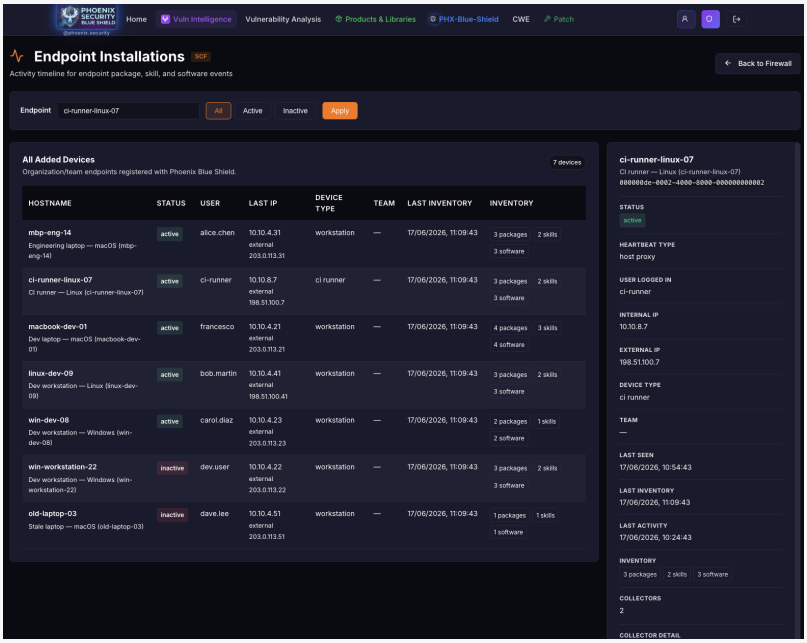

Blue Shield Endpoint Management

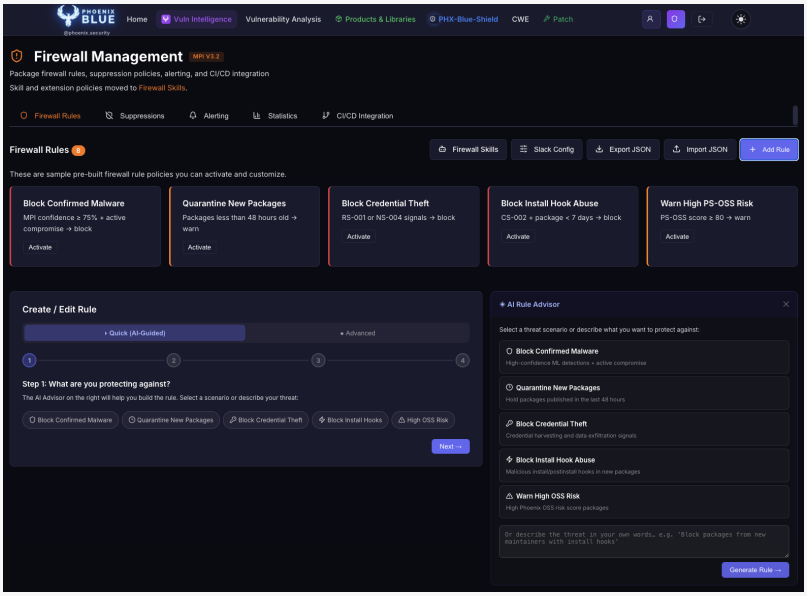

Blue Shield Package Rules

The structural truth

Across all 59 campaigns in Phoenix's corpus, zero CVEs were assigned during active exploitation. These attacks are not code defects — they are abuses of trust: a maintained package, a trusted publisher, a legitimate CI identity. There is no flawed code path to assign a CVE to, and no advisory to wait for. A defense that begins with “wait for the advisory” has already lost.

And now the attacker has a reason

The mythos-level shift is that frontier models compress the discovery-to-exploit cycle toward zero. Generation has become cheaper than interpretation: AI can find and weaponize a flaw in hours, while defensive review still assumes a human committing once a day. A handful of organizations have privileged access to the most capable frontier systems; everyone else defends on a limited budget against commodity models that are catching up fast. When attackers have reasoning, defenders need it too — but pointed at curated context, not the whole internet, or the economics never work.

The solution: Blue Shield
Blue Shield is the answer to a threat that does not announce itself. It does not wait for a CVE. It decides what a package or skill actually does — behaviourally — and blocks it at the moment of install, everywhere an install can happen.
Behavioral, not a blocklist

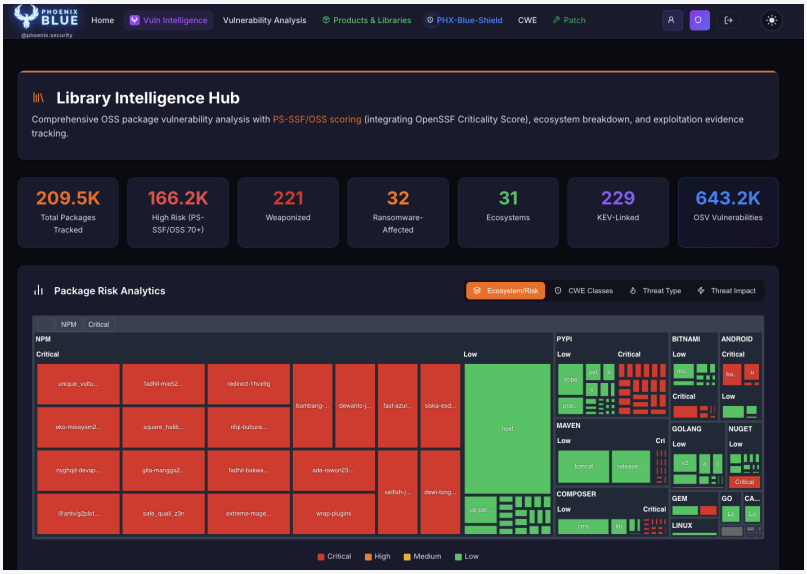

Phoenix Security Malware Alerts

A blocklist only knows what is already known to be bad. Blue Shield reads behavior: does this package reach for credentials, call out to infrastructure tied to a known campaign, drop an install hook, or carry a payload that fires on startup? Findings map to MITRE ATT&CK, and the verdict is explainable rather than a single opaque score. This is what lets it catch a fresh package from a campaign it has seen before, the day it is published, with no advisory in existence.
One verdict, every layer a package or skill enters
There is no single chokepoint that works. An install inside an agent skips your CI. A manual install on a laptop skips the agent. A poisoned skill never touches a package file. Blue Shield places a check at each layer and feeds them all from one intelligence backbone, Phoenix Blue, so a verdict made in one place holds in the others.

What Blue Shield does:
AI agent - Checks every install the agent proposes before it runs. Tells the agent which version is safe and which is not, or blocks it — and routes genuine cases to a human. Critical when agents run autonomously or in a crowd.
Developer workstation - Protects against malicious packages and compromised skills installed locally, including installs that never involve an agent.
CI/CD pipeline - Blocks malicious packages as they move through the build, or decorates the pull request to flag which package should and should not be used.

Blue Shield monitors across the package managers and systems attackers actually use — npm, PyPI, Maven, plus GitHub and Jenkins — and refreshes its intelligence continuously, so each verdict reflects where a package and its risk stand right now.

Endpoint monitoring and heartbeat
Blocking is half the job; you also need to know what you are running and whether your protection is alive. Blue Shield keeps a live inventory of every endpoint under its watch and a heartbeat from each collector, so the platform is the source of truth for what is installed and what is protected.

See every endpoint — workstations, CI runners, and agent sessions in one view, with which are active and which have gone quiet.

Know what's on them — the packages, agent skills, and IDE or VS Code extensions per endpoint, with the risky ones flagged.
Heartbeat and health — each collector reports in, so you can tell whether an endpoint is protected or has drifted. Each endpoint resolves to a single stable identity, so the workstation and agent collectors on the same machine appear as a single endpoint, not two.
The human stays in control

When an agent works on its own, someone still needs to be able to step in. Blue Shield gives the agent a clear signal and routes real decisions to a person to approve, with the package's behavior and campaign context attached. The AI assists the decision. It never replaces the engineer, and every block is reviewable.
Free for the core tier, today

Phil Moroni
Phoenix Security
+1 919-594-8888
email us here
Visit us on social media:
LinkedIn
Instagram
Facebook
YouTube
X
