TripleKey

FOR CYBER INSURANCE CARRIERS

The Carrier Underwriting Brief

Underwriting software supply chain risk with code level precision.

Cyber policies are priced on questionnaires. Losses happen in code. This brief maps the gap between the two, the data that closes it, and how carriers use continuous code level evidence to price, monitor, and defend cyber risk across the entire book.

For chief underwriting officers, heads of cyber, portfolio managers, actuaries, and claims leadership at carriers, MGAs, and reinsurers.

JUNE 2026

EXECUTIVE SUMMARY

Why this brief, and why now

Cyber carriers are exposed to a class of risk that traditional underwriting inputs were never designed to see. A SOC 2 attestation, a self attested SIG Lite, or an ISO 27001 certificate describes a single moment in time, and most high visibility breaches of the last three years happened to insureds who held all three. The signals that actually predict loss live inside an insured's software supply chain, and they change every day.

The numbers tell the story. Third party involvement in breaches doubled in a single year, from 15 percent to 30 percent of all breaches (Verizon DBIR 2025). Supply chain compromise costs $4.91 million per incident and takes 267 days to identify and contain, the longest of any attack vector (IBM Cost of a Data Breach 2025). New CVE volume crossed 48,000 in 2025, and FIRST projects a median of roughly 59,000 in 2026, the first year ever likely to pass 50,000 (NVD; FIRST 2026 Vulnerability Forecast). Meanwhile the U.S. cyber loss ratio reached 48.8 percent in 2024, up more than seven points year over year (AM Best).

This brief is built for carrier leadership. It maps the three structural gaps in cyber underwriting today, what the 2025 breach record means for the book, the new exposure AI generated code is introducing, the continuous data feed TripleKey delivers, and how a sample portfolio pilot works, from read only connection to live signal in under a week.

THE THESIS IN THREE LINES

- **Questionnaires cannot price a risk that changes daily.**
  The inputs cyber underwriting relies on expire on submission. Underwriters need a continuous signal, not an annual snapshot.
- **Aggregation is the book level threat.**
  One vulnerable open source library can expose hundreds of insureds at the same moment. Without an SBOM lens across the portfolio, correlated exposure surfaces only when claims arrive.
- **Evidence wins claims.**
  A daily, time stamped record of every insured's software state turns causation, coverage decisions, and subrogation from reconstruction into documentation.

THE UNDERWRITING GAP

Three places the current model leaks loss

01 Questionnaires expire on submission

Point in time attestations describe the day they were issued, not the day the breach happens. Roughly 60 percent of breaches involved exploiting a known vulnerability for which a patch was already available (Verizon DBIR 2025). That is a posture decay problem, and no annual artifact can detect it.

02 Aggregation risk is invisible

In September 2025, CISA issued an emergency alert over Shai-Hulud, a self replicating npm worm that compromised more than 500 packages by stealing developer credentials. A second wave in November 2025 backdoored roughly 796 packages carrying more than 20 million combined weekly downloads (CISA; Cloud Security Alliance). For a carrier, that is one event correlated across every insured that ships software, and without a portfolio wide SBOM view there is no way to quantify the accumulation until the claims arrive.

03 Claims defense lacks evidence

When a breach notice comes in, carriers need forensic certainty about what software was running, which dependencies were in place, and when warning signs first appeared. That record usually does not exist, and supply chain breaches already take 267 days to identify and contain, the longest of any vector (IBM 2025). Evidence assembled after the fact is expensive and contestable.

30%

of breaches now involve a third party, double the prior year

VERIZON DBIR 2025

$4.91M

cost per supply chain incident, 267 days to contain

IBM 2025

48,185

new CVEs published in 2025, roughly 132 per day

NVD / NIST

15 to 20%

of incoming CVEs now fully enriched under the NVD triage model

NIST, APRIL 2026

BREACH REPORT · 2025 DATA

What the 2025 breach record tells underwriters

The three annual references that anchor cyber loss data, the Verizon DBIR, IBM's Cost of a Data Breach study with Ponemon, and the Identity Theft Resource Center, all point the same direction: breach costs are concentrating in the United States, and third party and supply chain events are multiplying faster than any other category.

$10.22M

average U.S. breach cost in 2025, an all time high

IBM 2025

1,251

entities hit by supply chain breaches in 2025, up from 660

ITRC 2025

44%

of confirmed breaches involved ransomware, up from 32%

VERIZON DBIR 2025

12,195

confirmed breaches analyzed across 139 countries

VERIZON DBIR 2025

- **Costs are diverging, and supply chain is the slow burn**
  The global average breach cost fell 9 percent to $4.44 million in 2025, the first decline in five years, while the U.S. average reached an all time high of $10.22 million. Supply chain compromise was the second most prevalent and second costliest vector at $4.91 million per incident, and at 267 days it took the longest of any vector to identify and contain (IBM 2025).
- **Supply chain events reach further every year**
  The ITRC counted 1,251 entities affected by supply chain breaches in 2025, nearly double the 660 recorded in 2024, with each attack reaching more downstream targets than in prior years. For a carrier, downstream reach is the loss multiplier: one upstream compromise arrives as many simultaneous claims across the book.
- **Healthcare remains the costliest line**
  Healthcare breaches averaged $7.42 million, the costliest sector for the fourteenth consecutive year (IBM 2025). Roughly 57 million individuals were affected by healthcare breaches reported to HHS in 2025 across more than 640 large incidents, more than 80 percent of them from hacking and IT incidents (HHS Office for Civil Rights).
- **The cascade precedent is already on the record**
  The 2024 Change Healthcare attack remains the clearest public example: a single vendor compromise that exposed data on well over 100 million individuals and disrupted providers nationwide. Every number above argues for pricing on observable software posture rather than attestations.

EMERGING RISK · AI GENERATED CODE

AI is changing what is inside your insureds' software

Roughly three quarters of developers report using or planning to use AI coding tools (Stack Overflow Developer Survey). AI generated code, and the dependencies it pulls in, are now entering insured codebases at scale, and nothing on a renewal application captures it. The neutral record on what that means is getting clearer.

01 AI assistants recommend packages that do not exist

A USENIX Security 2025 study tested 16 code generating models across 576,000 code samples. Roughly one in five recommended packages did not exist, producing more than 205,000 unique fake package names. Commercial models hallucinated about 5.2 percent of the time, open source models about 21.7 percent (USENIX Security 2025).

02 The hallucinations are predictable, which makes them exploitable

When a prompt produced a fake package name, 43 percent of those names recurred across repeated queries. Attackers can register the names AI tools reliably invent and wait for developers to install them. The Python Software Foundation named the pattern slopsquatting. It joins typosquatting and dependency confusion as a standing registry threat.

03 AI driven attacks now show up in loss data

About 1 in 6 breaches in 2025 involved an AI driven attack, averaging $4.49 million per incident (IBM 2025). And the Shai-Hulud npm worm demonstrated that malicious dependencies can spread autonomously through the same package ecosystems AI tools draw from (CISA, September 2025).

THREE QUESTIONS THIS RAISES ON EVERY APPLICATION

- Does the insured know how much of its shipped code is AI generated, and can it say so under a warranty?
- Can the insured enumerate its dependencies today, not as of its last audit, and would a hallucinated or trojanized package show up?
- Would the insured detect a malicious package in its tree within 24 hours? A daily SBOM answers all three with evidence instead of attestation.

THE DATA CARRIERS RECEIVE

Live signals, from quote to claim

TripleScan runs a daily forensic scan of each insured's codebase and dependencies through a read only repository token. No agents, no pipeline changes, no engineering lift on the insured.

Tech Risk Score (0 to 100)

A composite score updated daily, trended over the policy period, and mapped to letter grades for executive reporting.

Live SBOM per insured

A full software bill of materials, refreshed daily and queryable across the entire book to surface aggregation exposure in real time.

CVE and license alerts

Severity ranked, mapped to affected insureds, with remediation status and patch latency tracked over time.

Contributor and provenance trail

Who wrote the code, where it came from, and what changed. Critical for IP exposure questions and post breach root cause work.

FOUR MOMENTS WHERE IT CHANGES THE ECONOMICS

- **Underwriting and quoting**
  Replace weeks of questionnaire back and forth with a score derived from the code the insured actually ships. Bind faster on better risks and tier premium and deductibles on observable hygiene.
- **Portfolio aggregation monitoring**
  The same day a critical CVE is disclosed in a widely used dependency, see exactly which insureds, which limits, and which lines of business are exposed.
- **Loss prevention and renewal**
  Configurable score thresholds trigger risk engineering outreach when an insured's posture decays, before a small problem becomes a covered loss, and arm tougher renewal conversations with evidence.
- **Claims investigation and defense**
  A time stamped SBOM and CVE history already exists for every insured. Use it to scope causation and defend coverage decisions.

HOW IT WORKS

From policy bound to live signal in under a week

TripleKey was designed to be a routine input to underwriting and risk engineering, not a heavy implementation. The insured grants a read only token. Everything else runs in the background, at no charge to the insured.

01

Insured opts in

At quote, at bind, or as a covered service in the policy.

02

Read only connection

A repository token is provisioned. No pipeline changes, no agents.

03

Daily forensic scan

TripleScan runs every 24 hours. Score, SBOM, CVE and license alerts land in the portal.

04

Carrier acts on signal

Underwriting prices, portfolio watches aggregation, claims has evidence.

UNDERWRITING INPUTS, COMPARED

CAPABILITY	QUESTIONNAIRES	ANNUAL SOC 2 / ISO	TRIPLEKEY
Refresh rate	Annual	Annual	**Daily**
Code level evidence	Not included	Sampled, point in time	**Full SBOM, every insured**
Aggregation visibility	Not possible	Not possible	**Cross portfolio queries**
Drift detection	None	None	**Score deltas, daily**
Forensic record for claims	If the insured kept one	Out of date by months	**Time stamped, immutable**
Lift on the insured	Weeks of questionnaires	Months of audit prep	**A read only token**

WHAT YOU RECEIVE DURING THE PILOT

Underwriting Brief

PER INSURED PDF

A one page summary written for an underwriter. Score, top exposures, renewal recommendations. Drops straight into the file.

Aggregation Reports

PORTFOLIO CVE ROLLUPS

The same day a critical CVE is disclosed, see exactly which insureds, which limits, and which lines of business are exposed.

Risk Engineering Pack

INSURED FACING REPORT

A non technical, executive ready report your risk engineers can send to the insured to drive remediation between renewals.

Claims Evidence

FORENSIC SBOM HISTORY

A time stamped record of every dependency, every score change, every alert, delivered to defense counsel as a single bundle.

WHY TRIPLEKEY

Built to put the security boundary outside the system it protects

TripleKey provides real time visibility into software risk and compliance. TripleScan, the company's scan engine, runs daily automated analysis of codebases and dependencies and produces the Tech Risk Score, SBOM, CVE alerts, license conflict detection, and contributor risk analysis. The company was founded in December 2023 and is headquartered in Brentwood, Tennessee. TripleKey monitors software risk in some of the most demanding regulated environments in the country, including enterprise healthcare, where it serves Community Health Network of Indianapolis. We work with cyber carriers, MGAs, and reinsurers across all industry verticals.

GRANTED USPTO INTELLECTUAL PROPERTY

In October 2025, the United States Patent and Trademark Office granted TripleKey US Patent 12,455,973 B1 for a novel approach to secure encryption. The patented system separates cryptographic key custody from the device performing the encryption work, requiring a removable hardware component to be physically present for any decryption to occur. This reflects TripleKey's broader architectural philosophy of moving security controls outside the system being protected, the same philosophy behind TripleScan's out of pipeline monitoring. The patent is wholly assigned to TripleKey and is enforceable through approximately January 2045.

SCALE WITH CONFIDENCE

Let's price the risk you can actually see.

Bring us a sample portfolio. We will show you what your insureds actually look like at the code level, and what that means for your loss ratio.

Talk to underwriting

triplekey.com · Response within 2 business days

TripleKey · 5205 Maryland Way, Suite 300, Brentwood, TN 37027 · triplekey.com