UNDERWRITING INTELLIGENCE FOR THE CYBER INSURANCE MARKET

Score the risk. Then improve it.

Rates are softening and loss ratios are climbing. Margin now lives in how you select the risk and how you improve it after bind. **HiveSec Engine for Insurance** does both, from a single evidence base.

01 · SELECT

BindSight™

An evidence-linked Cyber Tier Score, kept current by continuous scanning. Every penalty decomposes to an observed signal: a score your underwriters can defend and your actuaries can take apart.

02 · IMPROVE

Loss Control Desk™

Analyst-reviewed findings released to insureds, chased to closure, and evidenced for renewal.

WHAT'S DRIVING THE SCORE · SERVER-COMPUTED PENALTIES

● Observed

−45	Known-exploited vulnerability (CISA KEV)	Vulnerability Management
−25	Ransomware-linked CVE (CISA KEV)	Vulnerability Management
−18	High exploit probability (EPSS ≥ 0.5)	Vulnerability Management
−10	Internet-exposed RDP	Network Exposure
−8	Confirmed-present findings	Vulnerability Management

Each figure is the point penalty the signal costs the Security Score (`weight × count, capped`). Drill down from score → signal → finding → host.

BINDSIGHT · SELECT THE RISK AT QUOTE

SCORING

Cyber Tier Score: transparent by construction

A 0–100 Security Score and Exposure Score for every insured, resolved to a matrix cell and an A–F grade. The full scoring model is documented: every weight, every cap, every input.

VALIDATION

Confirmed-present findings

BindSight confirms each vulnerability is present and valid before it surfaces as a finding: HiveSec's AI agents test the exposure and peer-review the conclusion. Nothing is inferred, and what reaches the queue an underwriter can put in front of a broker.

THREAT INTELLIGENCE

Signals weighted by what drives claims

A known-exploited or ransomware-linked CVE costs an insured more points than a routine finding. Weightings draw on CISA KEV and EPSS, alongside exposure signals: internet-facing RDP, exposed databases, SMB/NetBIOS, and cleartext legacy protocols.

WORKFLOW

Underwriting queue & renewal triage

A queue ordered by bind-by date. Renewals due in the next 30 days, accounts awaiting a score, and grade drift across the book: the desk's priorities, already in order.

QUOTE CONDITIONS

What-if remediation projection

Tick the exposures an insured could fix and project the new score: a ready-made subjectivity list at quote and a loss-control agenda mid-term.

PORTFOLIO

Your whole book, continuously monitored

Score-over-time for every insured, cohort percentiles (p10–p90), grade distribution, and priority actions ranked across the whole book. When an aggregation event breaks, answer "which of our insureds are exposed?" from live observations.

EXPOSURE × SECURITY · THE APPETITE MAP

SECURITY · STRONG → WEAK ↓ · EXPOSURE →

Two axes, one decision

Inherent exposure of the attack surface, plotted against demonstrated security posture. Each insured lands in a matrix cell and grade that map to appetite, tiering, and rate. The book distribution shows where your portfolio concentrates.

LOSS CONTROL DESK · IMPROVE THE RISK IN-TERM

Selection protects the margin once; loss control protects it all year. The Loss Control Desk is how you reach out to policyholders mid-term: it takes confirmed findings to the insured, chases them to remediation, and carries the evidence into renewal. Improved posture earns its rate credit on the evidence, not on a questionnaire.

DETECTED
observed on scan → **ANALYST REVIEW**
validated, curated → **RELEASED TO INSURED**
briefing + fix steps → **AWAITING INSURED**
remediation tracked →

CLOSED
renewal evidence

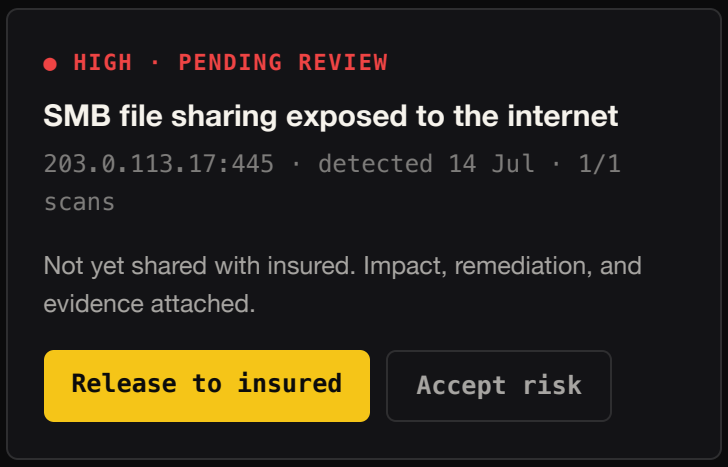

- **Analyst-reviewed release.** Every finding is reviewed before it reaches the insured: released or risk-accepted, finding by finding. The Engine confirms the finding is real; the analyst decides what the insured should see.
- **Briefings & case threads.** Send the insured a remediation briefing with evidence, impact, and fix steps attached; internal notes stay internal. Downloadable reports per case.
- **Book-wide outreach.** When a new critical exposure lands, reach every affected insured at once, with the fix and the evidence, before it becomes a claim.
- **Awaiting insured.** One queue for what the desk owes and what the insured owes; nothing stalls silently until renewal.
- **Renewal evidence.** The engagement trail (released, remediated, re-verified on scan) lands next to the score at re-rating.

> Run the desk with your own risk engineers, or have HiveSec's offensive-security analysts operate it as a managed service.

WHO IT'S FOR

ROLE	WHAT THEY GET

Underwriters	Defensible risk selection, faster submission triage, subjectivities with evidence attached.
Portfolio & exposure managers	Book distribution, D/F drift flags, overnight re-scoring of the in-force book.
Actuaries	A fully documented scoring model with auditable inputs for rating factors.
MGAs & MGUs	Show capacity providers why each risk was selected; API-first for fast-flow SME decisioning.
Risk engineers & loss control	The desk as their working queue: release decisions, briefings, and remediation tracked to closure.

HOW IT WORKS

01

Profile

Point the Engine at an account; it maps the insured's internet-facing estate.

02

Observe & validate

The Engine measures the signals, confirms each finding present, and applies threat-intelligence enrichment.

03

Score & steer

Cyber Tier Score, grade, matrix cell, and book percentile land in your queue and over the API.

04

Engage & improve

Brief the insured, track the fix, and carry the evidence into renewal.

See your book through the evidence.

Web platform and API · a workspace for each book or programme.

Book a demo → hello@hivesec.com

About HiveSec. HiveSec builds offensive-security technology. Our AI agents discover, test and confirm exposures at a scale manual review cannot reach. HiveSec Engine for Insurance puts that evidence to work: BindSight for underwriting, the Loss Control Desk for in-term engagement. HiveSec LTD, registered in England & Wales · hivesec.com

BindSight™ and Loss Control Desk™ are trade marks of HiveSec LTD (applications pending). Scoring model v1. Scores are evidence summaries of observed external signals; they are not a guarantee of insurability or loss outcome.