PART OF THE KAI AUTONOMOUS DEFENSE PLATFORM

Detection Intelligence

Detection in minutes. Alerts you can trust.

Not everything can be patched at once. Some fixes wait on a maintenance window, a vendor patch or a release, and the exposure stays open the whole time. Kai closes that gap with detection built for your environment.

74% to 12% reduction in false positive rate in a production security operations center

70+ detection rules authored in two hours during a single engagement

52% of CISOs name trust in automated decisions as the top barrier to more automation

Reported by 500 chief information security officers at enterprises above $500 million in revenue.

01 · THE PROBLEM

More rules is not more coverage

A feed arrives, a technique looks relevant, somebody writes a rule and it ships. Over the years the rule set grows into the hundreds and coverage is reported as a count.

A rule for a technique that cannot run in your environment is noise with a name. The exposures you genuinely have often have nothing watching them at all.

02 · THE SOLUTION

Detections written for the exposures you have

Kai writes, validates and maintains the rules that watch your real exposures, in the SIEM and EDR you already run.

03 · USE CASES

Autonomous detection engineering
Rules converted to your SIEM's language, validated against your schema and tested in several stages before they ship.
Coverage scales without headcount.

Zero day and supply chain threat detection
Threat intel matched to the software you run and turned into validated queries.
Defense before a CVE is published.

SIEM log optimization
Event types tiered Hot, Warm or Cold by relevance to live rules.
Lower SIEM cost, same coverage.

Every rule in Kai carries the reason it exists

150 RULES · 149 ACTIVE · 36 TECHNIQUES COVERED

RULE	INFORMED BY	TYPE	STATUS	TRUE POSITIVES
Privilege escalation via sudo exploit	Threat intel	Composite	ACTIVE	**90%**
Anomalous SSH and replication traffic on backup appliances	Confirmed vulnerability	Anomaly	RECOMMENDED	New
SMBv1 lateral movement	No SMBv1 hosts in scope		RETIRED	Cannot run here

Detection rules in Kai: what informed each one, whether it can run in your environment, and how often it earns its alert. Sample data.

04 · HOW KAI DOES IT

From a confirmed exposure to a rule that earns its alert.

1. **Intel read continuously**
Advisories, exploit activity and threat reporting, held until applicability is proven.

2. **Only rules that can run here**
Software, configuration and telemetry must all be present, or no rule is written.

3. **Written in your query language**
Authored for the platform it runs on, against fields you collect.

4. **Validated before shipping**
Fires on what it is scoped for, ignores the rest, adds no unmanageable volume.

5. **Tuned or retired**
Measured once live. Rules that stop earning their alerts are tuned or retired.

05 · OUTCOMES

Minutes

from confirmed exposure to a tuned rule running in your platform

74% to 12%

false positive rate, measured in a production security operations center

One record

so detection is scoped to the same assets and exposures as everything else

What you can trust, and what you can hand your CISO

Every change carries its evidence
The rollback path and the proof that justified it are recorded with the change, not in somebody's memory.

Every removal carries a reason
A finding cut from the queue carries a written reason you can defend in an audit.

The boundary is enforced
You set where Kai may act, per environment, asset class and change type. Outside it, nothing moves without a person.

Nothing deployed to start
Agentless and read only against the sources you already run. SOC 2 Type II. Nothing is replaced.

06 · WHAT KAI DOES

What Kai does that a threat feed cannot

01 Read intelligence continuously.
Advisories, exploit activity and threat reports processed as they arrive, held until applicable.

02 Test applicability first.
Software, configuration and telemetry checked. If any is missing, no rule is written.

03 Write the rule natively.
In your query language, against fields you actually collect, with the logic explained.

04 Validate before shipping.
Several stages of validation: fires when it should, stays quiet when it should not, volume known.

05 Measure and retire.
Alert frequency and true positive rate tracked. Rules leave the set as well as enter it.

06 Link coverage to exposure.
Written against exposures validated on your assets, as the interim control while a fix waits.

> Detection engineering does not scale. It is not just the cost, it is the complexity. No number of engineers can achieve what Kai does. It handles the hard stuff: behavioral signals, composite detections, automated tuning, so my team can focus on strategy.
>
> SECURITY LEADER, KAI CUSTOMER

07 · IN PRACTICE

Detection stays with Security

Kai resolves the exposure, the asset and the owner once. The detection is written from that record and delivered into the SIEM and EDR Security already runs. IT and Engineering see only the change.

Detection Intelligence

SIEM + EDR

A rule that earns its alert

01	02	03	04	05
Intel in	**Applicability**	**Rule written**	**Validated**	**Live, measured**
Advisories, exploits, threat reporting	Software, config and telemetry checked	In your SIEM or EDR query language	Multiple stages, alert volume known	Retired when it stops earning its alerts

APPLICABILITY, PER TECHNIQUE

✓ **Software present**	tomcat 9.0.110 on 2,412 hosts
✓ **Configuration permits**	AJP connector enabled
✓ **Telemetry exists**	auth + process events in SIEM
● **Already covered**	no existing rule matches

WHAT GOES INTO ONE RULE

TECHNIQUE	T1190, public-facing app exploit
SIGNALS	web, auth and process events joined
LOGIC	4 stage sequence per host, 5 min
SCOPE	only the 2,412 exposed hosts

Validated in 3 stages before it ships

70+
rules authored in two hours

74% → 12%
false positive rate, production SOC

Minutes
to coverage, inside your SIEM

Rules deploy only where the exposure is open, and retire when the fix lands.

Name an exposure. We will write the rule.
Pick a confirmed exposure in your estate. We will author the detection for the telemetry you collect, tell you the alert volume to expect, and deliver it into the platform you run.
Written in your query language. Tuned before it ships. Alert volume known in advance. kai.security

Get a demo