gruve

SOLUTION BRIEF

Threat hunting

Proactive discovery for threats that have slipped past the alerting layer

01

Business problem

Security controls are good at surfacing what they are tuned to catch. Threat hunting exists because sophisticated adversaries do not always behave in ways that line up neatly with existing detections, and because real compromise often spans identity, cloud, SaaS, endpoint, and administrative behavior before it becomes obvious enough to trigger an alert. Organizations can therefore have strong tooling and still miss the specific question that matters most: is an attacker already operating inside the environment in a way our controls have not meaningfully surfaced?

The need for threat hunting is reinforced by frontline data. Google Cloud M-Trends 2025 reports that exploits were the top initial infection vector in Mandiant's 2024 investigations and that global median dwell time rose to 11 days. Palo Alto Networks Unit 42 reports that the fastest intrusions it investigated in 2025 moved from initial access to data exfiltration in just 72 minutes, while Google Threat Intelligence Group reports adversaries using AI across targeting, research, phishing, and other parts of the attack chain. In other words, attackers can both linger quietly and move at machine-like speed when conditions favor them (Google Cloud / Mandiant; Palo Alto Networks Unit 42; Google Threat Intelligence Group).

Most teams know they should hunt more, but they rarely have the time, threat intelligence context, and cross-domain analytical bandwidth to do it in a disciplined way. The result is a common gap between "we have telemetry" and "we know whether an adversary is hiding in the parts of the environment that matter most right now." Threat Hunting closes that gap.

11 days

global median dwell time

72 minutes

attack timeline from initial access to data exfiltration

02

Why now

Attack patterns are becoming harder to predefine and easier to parallelize. IBM X-Force reported a 44% increase in attacks beginning with exploitation of public-facing applications in 2025, and Unit 42s 2026 reporting describes AI-assisted workflows that allow attackers to run reconnaissance and initial access attempts across hundreds of targets in parallel. Waiting for standard detections to do all the work is increasingly risky in high-value environments (IBM; Palo Alto Networks Unit 42).

At the same time, organizational change constantly creates new hunting questions: newly disclosed vulnerabilities, identity-platform shifts, SaaS expansion, remote access changes, executive onboarding, supply-chain dependencies, and changes in logging maturity all create reasons to hunt for specific behaviors before they become incidents. Threat hunting is therefore not just a “nice to have” for mature SOCs; it is part of how resilient organizations keep pace with attacker adaptation.

The AI lens matters here too, but it has to be framed correctly. Attackers are using AI to accelerate pieces of the intrusion lifecycle. Defenders can and should use AI to accelerate clustering, prioritization, and hypothesis development. But hunts still require human judgment to decide what is actually suspicious, what is explained by business context, and what evidence is strong enough to support a conclusion.

44%

increase in attacks beginning with exploitation of public-facing applications

03

Solution overview

Gruve's Threat Hunting engagement applies threat-informed, hypothesis-driven hunting techniques to uncover malicious or suspicious activity that may not have triggered existing controls. Hunts are built around the clients current risk picture - for example, a newly exploited technology, an adversary cluster relevant to the industry, a privileged-user population, a high-value SaaS environment, or a pattern of concern that deserves targeted follow-up across multiple evidence sources.

This is not an outsourced alert queue and not a generic one-size-fits-all hunting package. Each engagement is scoped around specific business and threat questions, then executed across the data sources most relevant to answering them. That usually means some combination of endpoint, identity, cloud, SaaS, email, and administrative activity, correlated with a threat-informed understanding of how attackers actually behave in the wild.

Gruve uses AI teammates to accelerate telemetry review, hypothesis support, and reporting preparation, but hunts remain investigator-led. Human analysts decide what to test, how to pivot, which anomalies are meaningful, and how to translate findings into defensible conclusions and operational improvements.

04

Gruve’s solution items

Gruve's solution items	Description
Threat-informed hunt hypothesis development	Builds the hunt around real adversary behavior, current client concerns, and the specific technologies or workflows most likely to reveal hidden compromise.
Multi-source behavior analysis	Correlates endpoint, identity, cloud, SaaS, and administrative evidence to test whether suspicious activity exists beyond what the alerting layer has already surfaced.
AI-teammate assisted clustering and pivoting	Uses AI to accelerate organization of large evidence sets and candidate patterns while keeping all meaningful hunt decisions and findings in human hands.
Detection and resilience gap analysis	Documents where existing detections, logging, access controls, or response workflows failed to surface behavior that warranted attention.
Executive and operational reporting	Delivers both the hunt conclusion and the concrete improvements needed to reduce future blind spots, whether or not active compromise is found.

Benefits

- **Finds what rules miss:** Tests for adversary behavior that may not cleanly match pre-existing alerts, signature logic, or workflow assumptions.
- **Reduces dwell time and blind spots:** Improves the organizations chances of finding threats earlier by looking deliberately in the places and patterns most likely to be missed.
- **Strengthens detections using real evidence:** Turns hunt results into improvements for logging, alerting, identity controls, and cross-domain visibility.
- **Prioritizes high-value questions:** Focuses scarce analyst attention on hypotheses tied to current threats, technology changes, and business-critical assets.
- **Improves executive confidence:** Gives leaders a clearer sense of whether hidden attacker activity has been actively tested for, rather than merely assumed absent.

05

Distinct service offerings

Tiers	Focus / scope	Core activities	Key deliverables	Target audience
Intelligence-led threat hunt	A hypothesis-driven hunt based on current adversary activity, threat intelligence, or industry-relevant campaign patterns.	Threat hypothesis design · telemetry review · adversary-pattern testing · reporting	Hunt findings report · observed behaviors · detection gaps · follow-on recommendations	CISOs · detection engineering · threat intelligence teams
High-value asset hunt	Focused hunting across privileged users, R&D systems, executive populations, crown-jewel SaaS platforms, or other high-impact targets.	Asset-centered analysis · privilege review · cross-source correlation · risk prioritization	High-value asset hunt brief · prioritized exposures · improvement actions	Security leadership · engineering leaders · executive risk owners
Post-change / post-exposure hunt	Targeted hunting after major platform changes, emerging CVEs, vendor notices, or control disruptions that may create short-term blind spots.	Change-aware hypothesis development · focused searches · anomaly validation · resilience recommendations	Targeted hunt summary · suspected activity analysis · resilience actions · monitoring plan	IR leaders · cloud teams · platform owners

06

Competitors

Competitor	Market focus / strengths	Limitations / Differentiation opportunity for Gruve
Unit 42 Managed Threat Hunting	Strong global brand with managed threat hunting tied to broad telemetry and frontline consulting.	A strong managed service option, but often optimized for continuous service operating models. Gruve is differentiated when the client wants a focused consulting engagement built around a custom hunt question and leadership narrative.
Red Canary	Recognized MDR provider with strong detection and response coverage across endpoint, identity, and cloud.	Excellent at ongoing detection and response, but less tailored to discrete threat-hunting engagements that must answer a specific executive or post-change question outside a standard MDR model.
Arctic Wolf	Broad security operations and managed hunting positioning with platform support and 24×7 operations.	Well suited to managed operations, but many clients still need a more bespoke hunt tied to a business event, a specific technology exposure, or a targeted adversary hypothesis.
Kroll	Experienced cyber responders with intelligence-led services and strong advisory posture.	Credible for broader cyber services, but Gruve differentiates with a tightly scoped, evidence-driven hunt approach that can bridge operational findings directly into leadership-ready recommendations.
Internal SOC / detection teams	Know the environment and can quickly query existing telemetry.	Internal teams often lack the dedicated time, threat-informed structure, and independent challenge function needed for disciplined hunts beyond day-to-day alert handling.

07

Gruve differentiators

- **Threat-informed and hypothesis-driven** rather than purely telemetry-driven or tool-driven.
- **Works across hybrid evidence sources,** including identity, SaaS, cloud, and endpoint context in a single hunt narrative.
- **Uses AI to accelerate hunt workflow** without allowing AI to decide what constitutes compromise.
- **Produces detection and resilience improvements** even when the hunt does not find active compromise.
- **Fits naturally** into post-exposure validation, executive assurance, or recurring resilience programs.

08

Case study

A multinational manufacturer asked Gruve to hunt for signs of adversary activity after a widely exploited edge technology vulnerability was disclosed and several peers in its sector reported follow-on identity abuse. The client had no declared incident, but leadership wanted to know whether the exposed technology had become a beachhead for hidden movement into identity systems, privileged SaaS access, or sensitive engineering environments.

Key results

- Identified suspicious identity and administrative behaviors that had not generated high-confidence internal alerts.
- Helped the client rule out broad destructive or ransomware activity while still acting on meaningful exposure indicators.
- Improved detection logic and monitoring priorities for the specific attack pattern relevant to the clients environment.
- Delivered a defensible threat-hunt narrative for leadership during a sector-wide risk event.

About the client

Large multinational manufacturer with a lean central security team, regionally distributed infrastructure, and heightened concern about supply-chain and edge-system exposure.

Challenges

The client did not need generic reassurance; it needed to know whether a specific exposure pattern had translated into follow-on activity in its environment. Existing detections were not designed for that question. They could flag known bad indicators, but they did not provide a structured way to test whether an attacker had used the technology exposure to pivot toward identity abuse, privilege escalation, or sensitive data access paths.

The volume of telemetry also worked against the client. There was enough data to investigate, but not enough dedicated hunting bandwidth to test multiple hypotheses quickly and consistently. Leadership needed an answer while the sector-wide concern was still current, not weeks later when the risk narrative had already moved on.

Solutions

- **Threat-informed hunt design:** Gruve built the hunt around the specific exposure scenario, the attacker behaviors most likely to follow it, and the clients highest-value identity and administrative surfaces.
- **Cross-domain hunting:** Investigators correlated edge, identity, SaaS, and endpoint context to test whether there was evidence of post-exposure foothold expansion that standard detections had missed.
- **AI-assisted clustering:** Gruve used AI teammates to accelerate organization of candidate patterns and reporting drafts, while investigators verified supporting evidence and judged significance manually.
- **Detection improvement guidance:** The engagement concluded not only with findings, but with concrete recommendations for logging, hunt cadence, and detection updates tied to the actual attack path the client cared about.

Results and benefits

- Gave leadership a more credible answer than either silence or generic reassurance would have provided.
- Improved detection coverage for the specific threat pattern facing the clients sector.
- Allowed the client to focus remediation on evidence-backed issues rather than broad, high-friction emergency measures.
- Demonstrated the practical value of consulting-led threat hunting even in organizations with mature telemetry estates.

Don’t wait for an alert to tell you everything is fine.

See how Gruve’s Threat Hunting engagement helps your team test for hidden adversary activity, reduce blind spots, and improve resilience with evidence-backed conclusions.

Website: https://www.gruve.ai/

Email: contact@gruve.ai